

# Cisco NDFC Fabric Controller Configuration Guide, Release 12.0.x

**First Published:** 2021-09-30

**Last Modified:** 2022-02-07

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

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

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

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

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

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

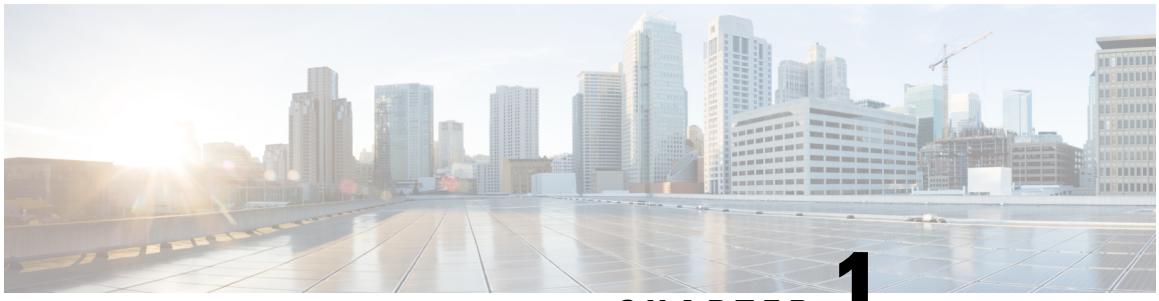
Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual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

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

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2021–2022 Cisco Systems, Inc. All rights reserved.



CHAPTER 1

# Overview


- Know your Web UI, on page 1
- User Feedback, on page 2
- Cohosting of NDFC Managed mode with Nexus Dashboard Insights, on page 3


## Know your Web UI

When you launch the Cisco Nexus Dashboard Fabric Controller Web UI for the first time, the **Feature Management** window opens. After you choose a deployment type, the left pane displays menu relevant to the personality.

The top pane displays the following UI elements:

- **Home** icon – Click to view One view on the Nexus Dashboard setup.
- **Nexus Dashboard** – Click to view One view on the Nexus Dashboard setup.
- **Feedback** – Allows you to provide feedback on Cisco Nexus Dashboard Fabric Controller. For instructions, see User Feedback, on page 2.
- **Help** – Click on **Help** to see a drop-down list with the following options:
  - **About Nexus Dashboard** – Displays the version of the Cisco Nexus Dashboard on which Cisco Nexus Dashboard Fabric Controller is deployed.
  - **Welcome Screen** – Displays What's New information. You can choose to see this page every time you launch the Web UI.
  - **Help Center** – Click to view the Help Center page. You can access various product documents from this page.

    Scroll to the end of the page to view the services installed on Nexus Dashboard. Click on the Service to view **Help Center**.
- **User Role** – Displays the role of the user who is currently logged in, for example, **admin**. Click on the username to see a drop-down list with the following options:
  - **User Preferences** – Allows you to view the Welcome screen on every login.
  - **Change Password** – Allows you to change the password for the current logged-in user.

    If you are a network administrator user, you can modify the passwords of other users.

- **Logout** – Allows you to terminate the Web UI and return to the login screen.

- **Cisco Nexus Dashboard Fabric Controller Persona** – Specifies the deployment persona – **Fabric Controller** or **SAN Controller** or **Fabric discovery**.
- **Alerts and Notifications** – You can view the alerts and event notifications by clicking the Alerts and Notifications icon, next to the **Help** icon, in the top pane of Cisco Nexus Dashboard Fabric Controller.
- **Alarms** – The **Alarms** icon flashes when there is an Alarm or when thresholds exceed for your Cisco Nexus Dashboard Fabric Controller Deployment. Click on the flashing **Alarms** icon to view the messages. The following alarms are displayed.
  - **Interfaces Limit Exceeded** – If the maximum number of endpoints across all fabrics exceed 100K, the Alarms icon flashes and displays a message.
  - **High Availability (HA) State** – The HA state notification appears when the Native HA setup is not synchronized; when one of the nodes or both nodes may have stopped, failed, or not ready; or when the Alarms icon flashes. If the HA setup is synchronized, the notification clears either in 30 minutes (during the polling cycle) or when you log out and log in to the Cisco Nexus Dashboard Fabric Controller Web UI.
  - **Application down** – If one or more applications are down, an error appears. An alarm message appears when the applications go online or offline.

General icons on UI:

- **Hamburger** icon – Click on **Hamburger** icon adjacent to product name on home screen to minimize the menu items on home screen or to view menu items in details.
- **Refresh** icon – Click refresh icon to refresh and load screen.

# User Feedback

Cisco Nexus Dashboard Fabric Controller allows you to provide feedback on the application. You can ask for new features/enhancements using this feature. The request is sent to Cisco Nexus Dashboard Fabric Controller marketing engineers, who evaluate the requirement and include the features or enhancements in the upcoming releases.

To provide feedback using the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Procedure

**Step 1** Click on **Feedback** on the right top corner of the Nexus Dashboard Fabric Controller application.

When using this for the first time, you must configure the DNS and proxy server on the Cisco Nexus Dashboard.

**Step 2** To establish a connection, go to **Cisco Nexus Dashboard** on your browser and perform the following steps:

a) On the Cisco Nexus Dashboard Web UI, choose **Infrastructure > Cluster Configuration**.

The **Cluster Configuration General** tab appears.

b) In the **Proxy Configuration** area, click **Edit** icon.
c) In the **Servers** area, click **Add Server**.
d) Choose the protocol type – **HTTP** or **HTTPS**.
e) In the **Server** field, enter the IP address.
f) Enter the **Username** and **Password** in the respective fields.
g) Click the tick icon to confirm. Click the wrong icon to delete.
h) In the **Ignore Hosts** area, click **Add Ignore Host**.
i) Enter the **Hostname** and click tick icon to confirm. Click wrong icon to delete.
j) Click **Save** to configure the proxy server.

**Note**
Wait for up to five minutes for the Proxy configuration to propagate to the Nexus Dashboard Fabric Controller application.

**Step 3** On the Cisco Nexus Dashboard Fabric Controller Web UI, click **Feedback**.

**Step 4** On the **Feedback** panel, click on the stars to let us know how you feel about Nexus Dashboard Fabric Controller.

**Step 5** In the **Make a suggestion** field, enter your suggestion/feedback.

**Step 6** If you choose Cisco to contact you regarding the feedback, check the **Cisco may contact me about my feedback** check box.

**Step 7** Provide your **Name** and **Email** in the respective fields.

# Cohosting of NDFC Managed mode with Nexus Dashboard Insights

From Release 12.1.1e, you can host NDFC Fabric Controller persona and Nexus Dashboard Insights on the same Nexus Dashboard Cluster in Managed mode to manage fabrics and Nexus Dashboard Insights to monitor the same fabrics. Note that NDFC in Fabric discovery mode, that is, monitored mode with NDI on the same Nexus Dashboard cluster is supported with NDFC Release 12.0.2f. Cohosting requires 4 physical Nexus Dashboard nodes for a maximum scale of up to 50 switches. This functionality is also supported on NDFC Release 12.1.1e with the corresponding paired Nexus Dashboard Insights release.



**Note** Nexus Dashboard deployed on KVM doesn't support cohosting NDFC and Insights service on the same Nexus Dashboard cluster.



**Note** For cohosting NDFC and Insights on the same Nexus Dashboard cluster, the Nexus Dashboard nodes must be Layer 2 adjacent. Support for Layer 3 adjacency for cohosting deployments will be introduced in future releases.

The following table shows the compatible versions for Nexus Dashboard and services.

| Services | Compatible Version |
|---|---|
| Nexus Dashboard | 2.1.1e2.1.2d |
| Nexus Dashboard Insights | |
| Nexus Dashboard Fabric Controller | 12.0.2f12.0.1a |

The following table shows the system requirements for Nexus Dashboard.

| Specification | Supported Scale |
|---|---|
| Number of physical Nexus Dashboard nodes | 5 |
| Number of switches supported | 50 |
| Number of flows supported in Nexus Dashboard Insights | 10000 |

### Installation of NDFC and NDI on the same Nexus Dashboard

Cisco NDFC can be cohosted with Nexus Dashboard Insights on the same Nexus Dashboard.

**Before you begin**

- Ensure that you've installed the required form factor of Cisco Nexus Dashboard. For instructions, refer to Cisco Nexus Dashboard Deployment GuideCisco Nexus Dashboard Deployment Guide.
- Ensure that you meet the requirements and guidelines described in *Prerequisites* section in *Cisco NDFC Installation Guide*.
- The Cisco DC App Center must be reachable from the Nexus Dashboard via the Management Network directly or using a proxy configuration. Nexus Dashboard proxy configuration is described in *Cisco Nexus Dashboard User GuideCisco Nexus Dashboard User Guide*.
- If you are unable to establish the connection to the DC App Center, skip this section and follow the steps described in *Installing Services Manually* section in *Cisco NDFC Installation Guide*.
- Ensure that the services are allocated with IP pool addresses on the Cisco Nexus Dashboard. For more information, refer to Cluster Configuration section in *Cisco Nexus Dashboard User GuideCisco Nexus Dashboard User Guide*.

### Installing Nexus Dashboard

Install the required form factor of Cisco Nexus Dashboard. For instructions, refer to Cisco Nexus Dashboard Deployment Guide.

### Installing NDFC

Refer to *Cisco NDFC Installation Guide*.

Configure NDFC sites on Nexus Dashboard. Refer to the *Adding Sites* section in the Cisco Nexus Dashboard Deployment GuideCisco Nexus Dashboard Deployment Guide.

### Installing NDI

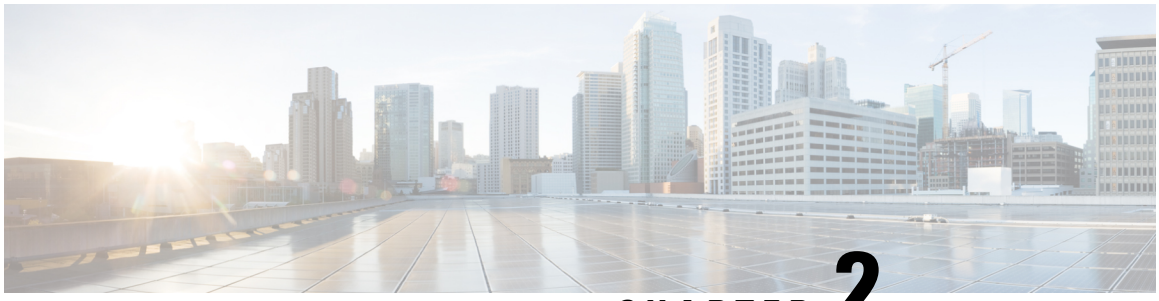
On the same Nexus Dashboard set up, install the Nexus Dashboard Insights service. Refer to Cisco Nexus Dashboard Insights Deployment Guide, for more information.

### Post Installation

After installing compatible versions of NDFC and NDI on the 5-node physical Nexus Dashboard, launch NDFC as Fabric (LAN) Controller. Create Fabric, discover and import switches on NDFC fabric. Nexus Dashboard automatically identifies the NDFC fabric and lists on the Sites page as entities.



**Note** You must provide the password for each of the sites in the Nexus Dashboard site manager.



CHAPTER 2

# Dashboard

The intent of the **Dashboard** is to enable network and storage administrators to focus on particular areas of concern around the health and performance of data center switching. This information is provided as 24-hour snapshots.

The functional view of LAN switching consists of seven dynamic dashlets that display information in the context of the selected scope by default.

The various scopes that are available on the Cisco Nexus Dashboard Fabric Controller Web UI are:


- Overview, on page 7
- Viewing vCenter VMs, on page 8
- Viewing Kubernetes Pods, on page 9
- Endpoint Locator Dashboard, on page 11


## Overview

From the left menu bar, choose **Dashboard** > **Overview**. The **Overview** window displays the default dashlets. The dashlets display donuts summary.

The following are the default dashlets that appear in the **Overview** dashboard window:

| Dashlet | Description |
|---|---|
| Fabric Health | Displays the fabric health status, and a number in the donut depicting total number of fabrics. The **Fabric Health** status is based on the severity of the highest outstanding alarm for the fabric or its member switches. |
| Events Analytics | Displays events with **Critical**, **Error**, and **Warning** severity. |
| Switches Configuration | Displays the switches inventory summary information such as the switch models and the corresponding count. |
| **Switches** | |
| Switch Health | Displays the switches health summary with the corresponding count. The **Switch Health** status is based on the severity of the highest outstanding alarm for the switch or its interfaces. |

| Dashlet | Description |
|---|---|
| Switch Roles | Displays the switches roles summary and the corresponding count. Displays the number of access, spine and leaf devices. |
| Switch Hardware Version | Displays the switches models and the corresponding count. |
| Switch Software Version | Displays the switches software version and the corresponding count. |
| **Performance Collector** | Displays the performance collection information.<br>• Click **Stop collector** to stop performance collection information.<br>• Click **Start collector** to restart the performance collection information. |
| **Reports** | Displays switch reports. |

# Viewing vCenter VMs

UI Path: **Dashboard** > **vCenter VMs**

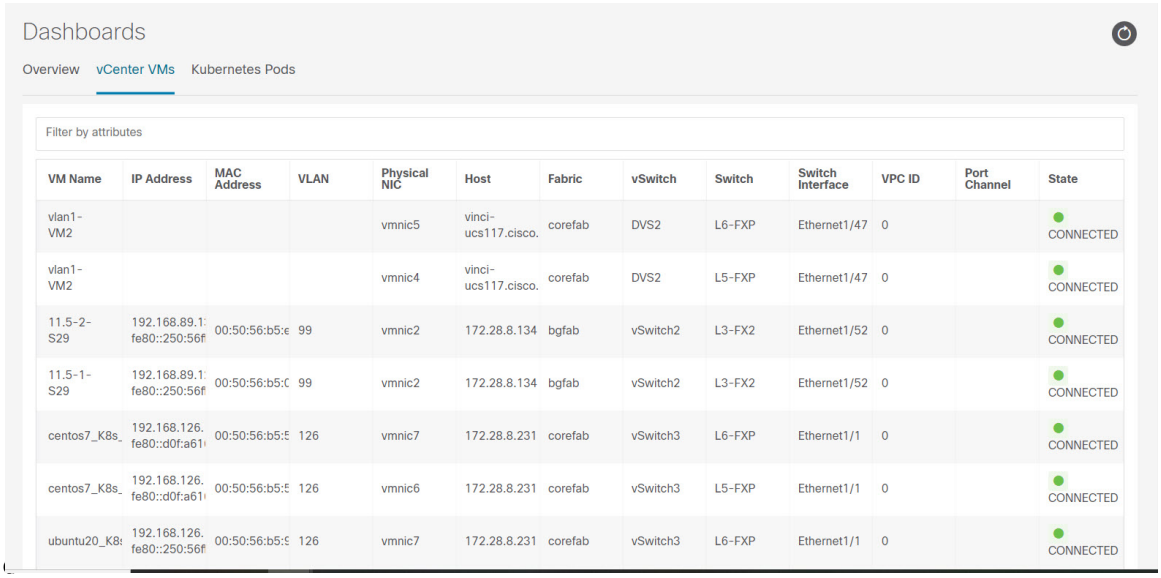


**Note** You can view the Virtual Machine details for the added vCenter cluster on dashboard and topology window. Navigate **Dashboard** > **vCenter VMs**.

The vCenter VMs tab displays the following details of VMs:

- VM Name, its IP address and MAC address
- Name of the compute where the VM is hosted
- Switch name that is connected to a VM, switch's IP address, MAC address, and interface
- Port channel ID and vPC ID (if connected to a VPC)
- VLAN VM configured on
- Power state of the VM
- Physical NIC of the Compute host

You can search and filter VMs by using **filter by attributes** search




To view VMs on Fabric window, navigate to **LAN** > **Fabrics**, double-click on required fabric. On **Fabric Overview** window, choose **Virtual Infrastructure** > **Virtual Machine VMs**.

To view VMs on Switch window, navigate to **LAN** > **Switches**, double click on required switch. On **Switch Overview** window, choose **Virtual Infrastructure** > **Virtual Machine VMs**.

# Viewing Kubernetes Pods

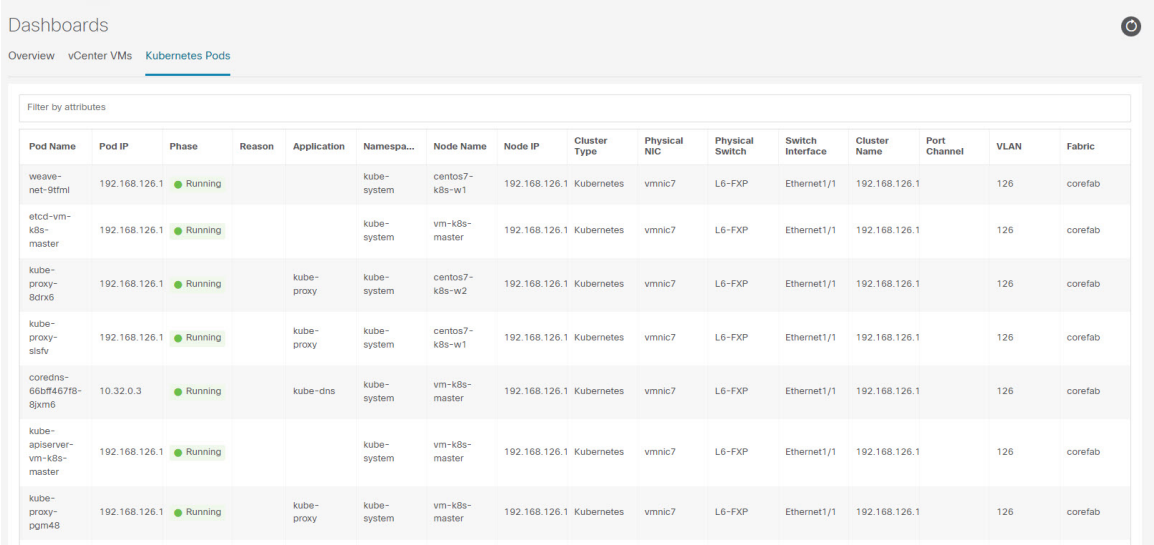
**Note** This is a preview feature in Nexus Dashboard Fabric Controller, Release 12.0.1a. We recommend that you use this feature marked as BETA in your lab setup only. Do not use these features in your production deployment.

UI Path: **Dashboard** > **Kubernetes Pods**

You can view Kubernetes pods on Fabrics window, navigate **LAN > Fabrics**, double-click on required fabric, it navigates to **Fabric Overview** window, click **Virtual Infrastructure** > **Kubernetes Pods**.

You can view Kubernetes pods on Switch window, navigate **LAN > Switches**, double-click on required switch, it navigates to **Switch Overview** window, click Virtual **Infrastructure > Kubernetes Pods.**

You can search and filter kubernetes pods by using **filter by attributes** search field.

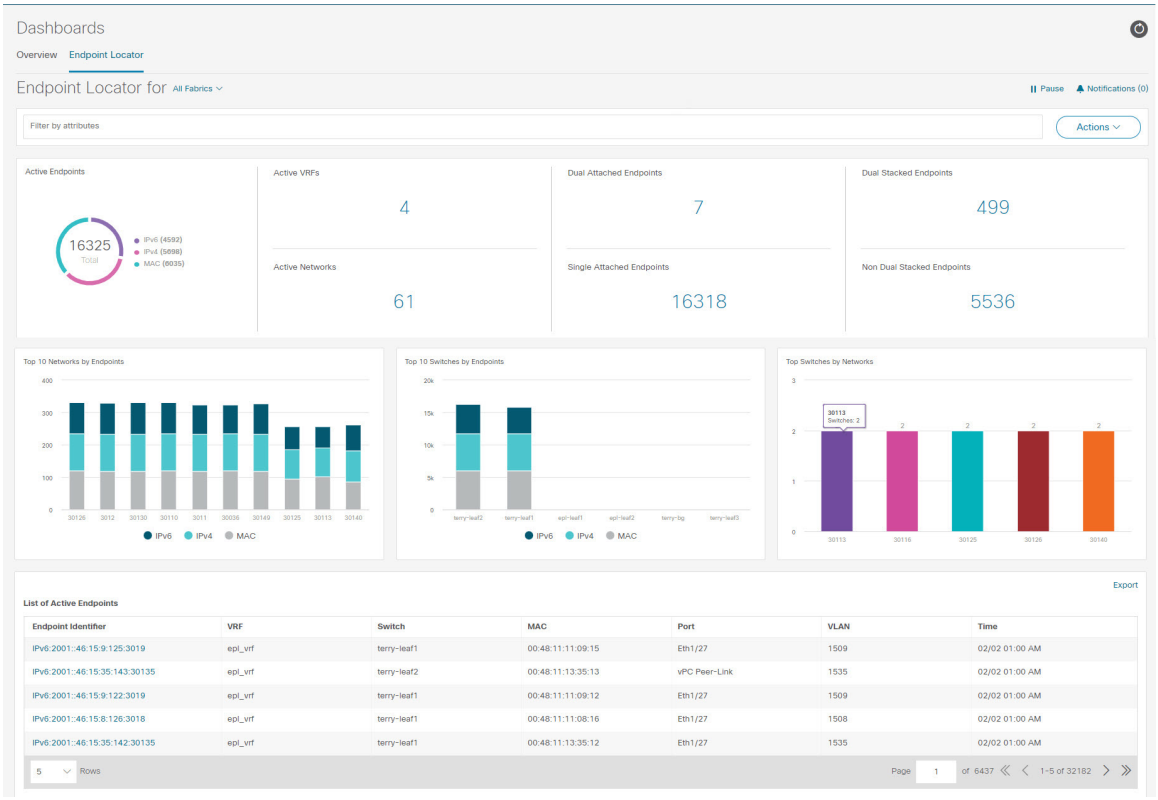



The following table describes the fields and description on the window.

| Field | Description |
|---|---|
| Pod Name | Specifies the name of the Kubernetes pod. |
| Pod IP | Displays the IP address of the Kubernetes pod. |
| Phase | Specifies the phase (state) of the pod. |
| Reason | Specifies the reason. |
| Applications | Specifies the applications of the pod. |
| Namespace | Specifies the namespace of the pod. |
| Node Name | Specifies the node name of the pod. |
| Node IP | Specifies the node IP address. |
| Cluster Type | Displays the type of cluster. |
| Physical NIC | Displays the physical NIC of the node. |
| Physical Switch | Specifies the physical switch connected to cluster node. |
| Switch Interface | Specifies the switch interface connected to cluster node. |
| Cluster Name | Specifies the name of the cluster. |
| Port Channel | Specifies the port channel (if cluster node is connected to a VPC). |
| VLAN | Specifies the VLAN. |
| Fabric | Specifies the fabric name. |

# Endpoint Locator Dashboard

To explore endpoint locator details from the Cisco Nexus Dashboard Fabric Controller Web UI, choose **Dashboard > Endpoint Locator**. The **Endpoint Locator** dashboard is displayed.




**Note** Due to an increase in scale, the system may take some time to collect endpoint data and display it on the dashboard. On bulk addition or removal of endpoints, the endpoint information displayed on the EPL dashboard takes a few minutes to refresh and display the latest endpoint data.

- You can initiate a search by using the available options in the **filter by attributres** search bar field.

  You can also filter and view the endpoint locator details for a specific **Switch, VRF, Network**, and **Type** by using the respective drop-down lists. You can select MAC type of endpoints as a filter attribute. The name of the network is also displayed in the **Network** drop-down list. By default, the selected option is **All** for these fields. You can display endpoint data for a specific device by entering the host IP address, MAC address, or the name of the virtual machine in the **Search Host IP/MAC/VM Name** field.

- You can click **All fabrics** drop-down list to view endpoint locator details for all fabrics or required fabric.

  An alarm is generated if there are any endpoint related anomalies. Click the **Pause** icon to temporarily stop the near real-time collection and display of data. By default **Run** is chosen. Click **Notification** icon to view the notification details.

- Click **Actions** > **Endpoint Search**. For more information, refer to Endpoint Search, on page 15.
- Click **Actions** > **Endpoint Life**. For more information, refer to Endpoint Life, on page 16.
- Click **Actions** > **Resync** to syncing to the data currently in the Route Reflector (RR). However, historical data is preserved. We recommend not clicking **Resync** multiple times as this is a compute-intense activity.

  In certain scenarios, the datapoint database may go out-of-sync and information, such as the number of endpoints, is not displayed correctly due to network issues such as:

  - Endpoint moves under the same switch between ports and the port information needs some time to be updated.
  - An orphan endpoint is attached to the second VPC switch and is no longer an orphan endpoint.
  - NX-API not enabled initially and then enabled at a later point in time.
  - NX-API failing initially due to misconfiguration.
  - Change in Route Reflector (RR).
  - Management IPs of the switches are updated.

- Click **Notifications** icon to display a list of the most recent notifications.

  The **Endpoint Locator Notifications** window appears.

  Information such as the time at which the notification was generated, the description of the notification, severity level is displayed.

  Notifications are generated for events such as duplicate IP addresses, duplicate MAC-Only addresses, VRF disappears from a fabric, all endpoints disappear from a switch, endpoint moves, endpoints on a fabric going to zero, when endpoints are attached to a switch, when a new VRF is detected, and when the RR BGP connectivity status changes. The RR connected status indicates that the Nexus Dashboard Fabric Controller can connect to the RR through BGP (Nexus Dashboard Fabric Controller and RR are BGP neighbors). The RR disconnected status indicates that the RR is disconnected and the underlying BGP is not functioning.

  You can initiate a search by using the available options in the **filter by attributres** search bar field.

The top pane of the window displays the following information:

The top pane of the window displays the number of active endpoints, active VRFs, active networks, dual attached endpoints, single attached endpoints and dual stacked endpoints, for the selected scope. Support for displaying the number of dual attached endpoints, single attached endpoints and dual stacked endpoints has been added. A dual attached endpoint is an endpoint that is behind at least two switches. A dual stacked endpoint is an endpoint that has at least one IPv4 address and one IPv6 address.

- Historical analysis of data is performed and a statement mentioning if any deviation has occurred or not over the previous day is displayed at the bottom of each tile.

  Click any tile in the top pane of the EPL dashboard to go to the Endpoint History window.

The 'middle pane' of the window displays the following information:

- **Top 10 Networks by Endpoints** - A pie chart is displayed depicting the top ten networks that have the most number of endpoints. Hover over the pie chart to display more information. Click on the required section to view the number of IPv4, IPv6, and MAC addresses.

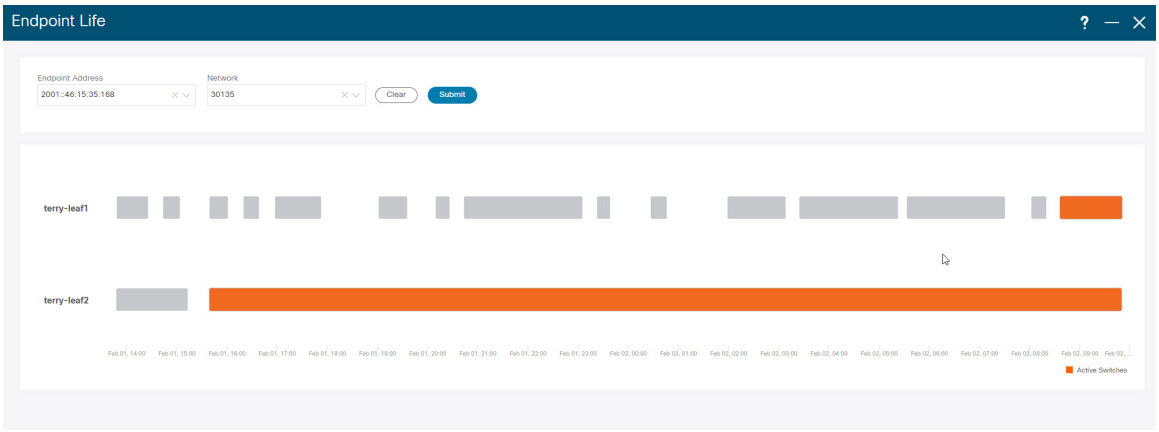
- **Top 10 Switches by Endpoints** - A pie chart is displayed depicting the top ten switches that are connected to the most number of endpoints. Hover over the pie chart to display more information. Click on the required section to view the number of IPv4, IPv6, and MAC addresses.
- **Top Switches by Networks** - Bar graphs are displayed depicting the number of switches that are associated with a particular network. For example, if a vPC pair of switches is associated with a network, the number of switches associated with the network is 2.

The 'bottom pane' of the window displays the list of active endpoints.

If a virtual machine has been configured, the name of the VM is displayed in the **Node Name** field. Note that it can take up to 15 minutes for the name of the VM to be reflected in the EPL dashboard. Until then, the EPL dashboard displays **No DATA** in the **Node Name** field.

Click **Export** to download the list of active endpoints in .csv format.

Click on required endpoint identifier, a slide-in pane appears and the related details are displayed. Click **Endpoint Life**. The **Endpoint Life** window appears for selected endpoint identifier. For more information, refer to Endpoint Life, on page 16.



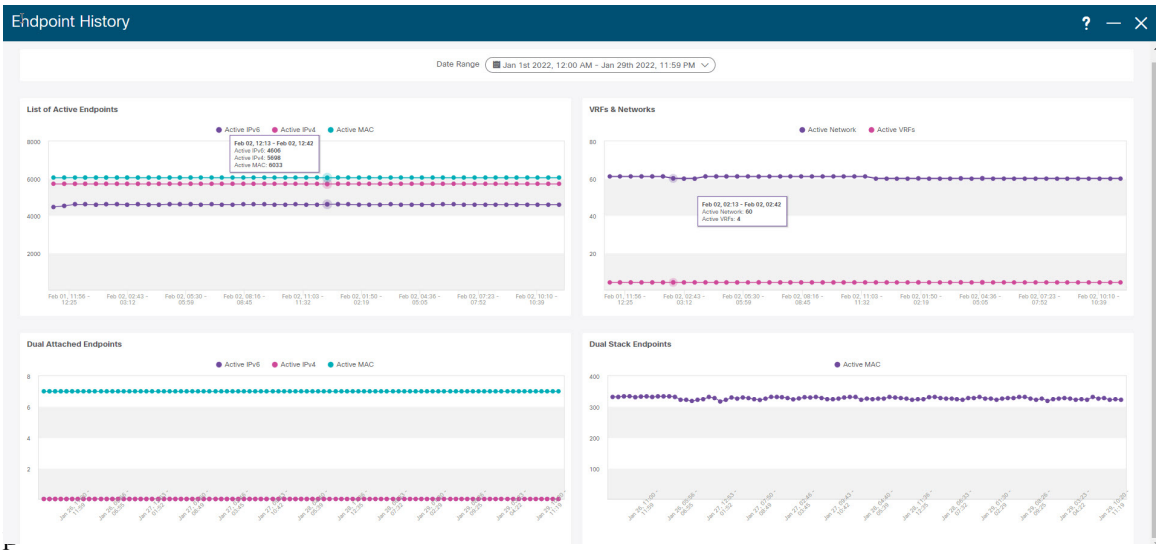

Click the search icon in the **Endpoint Identifier** column to search for specific IP addresses.

Consider a scenario in which EPL is first enabled and the **Process MAC-Only Advertisements** checkbox is selected. Then, EPL is disabled and enabled again without selecting the **Process MAC-Only Advertisements** checkbox. As the cache data in elasticsearch is not deleted on disabling of EPL, the MAC endpoint information is still displayed in the EPL dashboard. The same behavior is observed when a Route-Reflector is disconnected. Depending on the scale, the endpoints are deleted from the EPL dashboard after some time. In certain cases, it may take up to 30 minutes to remove the older MAC-only endpoints. However, to display the latest endpoint data, you can click **Resync**.

# Endpoint History

Click any tile in the top pane of the EPL dashboard to go to the **Endpoint History** window. A graph depicting the number of active endpoints, VRFs and networks, dual attached endpoints and dual stacked MAC endpoints at various points in time is displayed. The graphs that are displayed here depict all the endpoints and not only the endpoints that are present in the selected fabric. Endpoint history information is available for the last 30

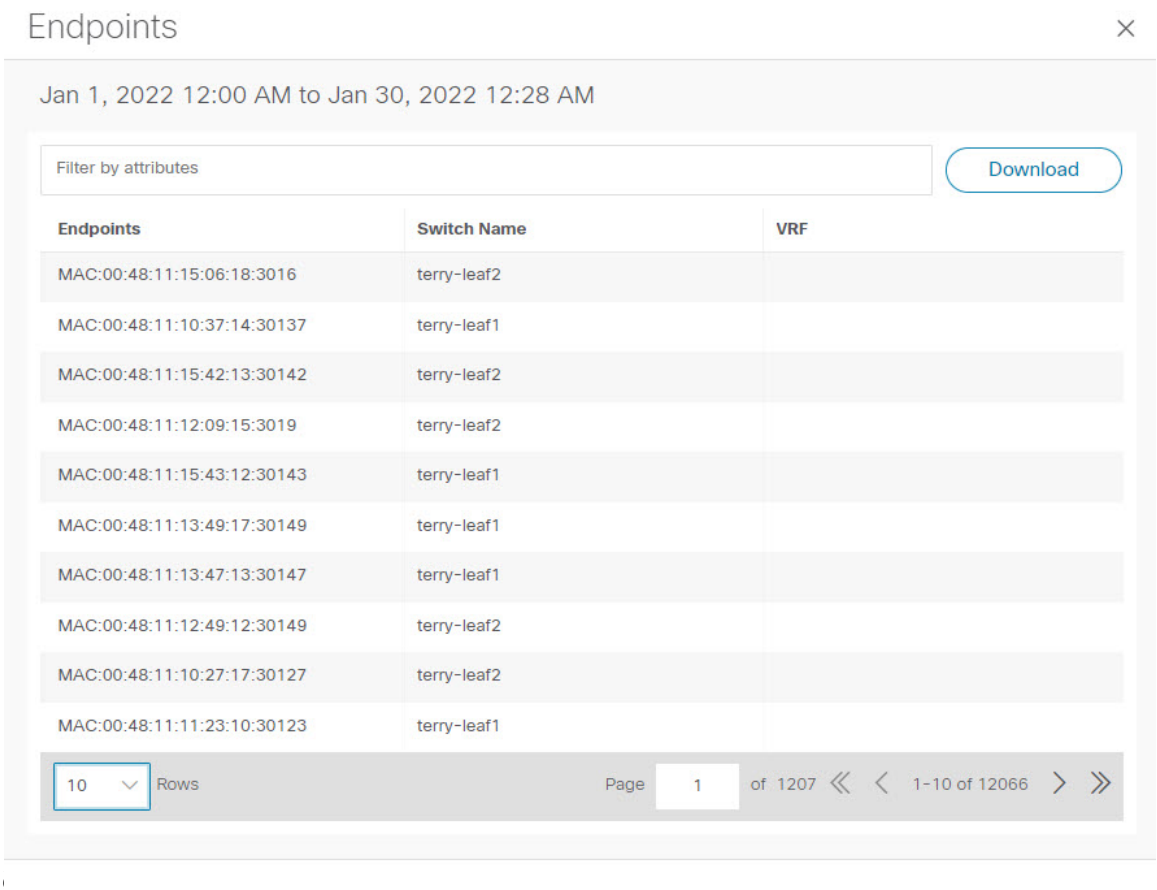
days amounting to a maximum of 100 GB storage




Hover over the graph at specific points to display more information. The points in the graph are plotted at 30-minute intervals. You can also display the graph for a specific requirement by clicking the color-coded points at the bottom of each graph. For example, click on all color-coded points other than **active (IPv4)** in the Active Endpoints window displayed above such that only **active (IPv4)** is highlighted and the other points are not highlighted. In such a scenario, only the active IPv4 endpoints are displayed on the graph. You can also click on the required color-coded points at the bottom of the graph to display the graph for a specific requirement. For example, hover over **active (IPv4)** to display only the active IPv4 endpoints on the graph.

Click on any point in the graph to display a window that has detailed information about that point of time. For example, click on a specific point in the **Active Endpoints** graph to display the **Endpoints** window. This window has information about the endpoints along with the name of the switch and the VRF associated with

the endpoint. Click **Download** to download the data as a CSV



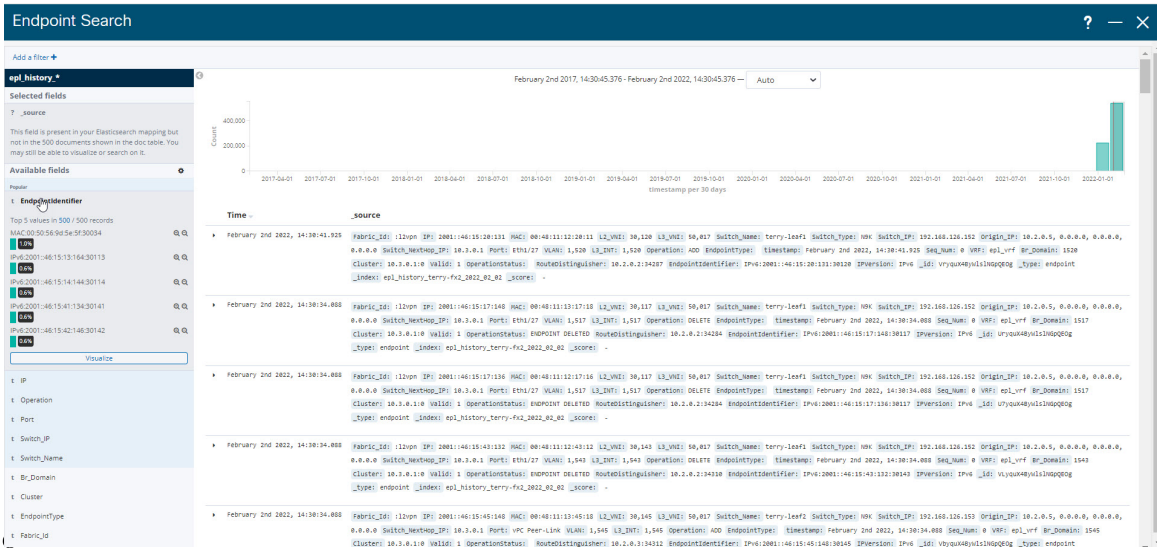

# Endpoint Search

UI Path: **Dashboard** > **Endpoint Locator**.

On **Endpoint Locator** window, click **Actions** > **Endpoint Search** to view a real-time plot displaying endpoint events for the period specified in a date range.

**Note** You cannot change time on the clock icon. Ignore the tooltip to change the time.

The results displayed here are dependent on the fields listed under **Selected fields** located in the menu on the left. You can add any field listed under **Available fields** to **Selected fields** to initiate a search using the required



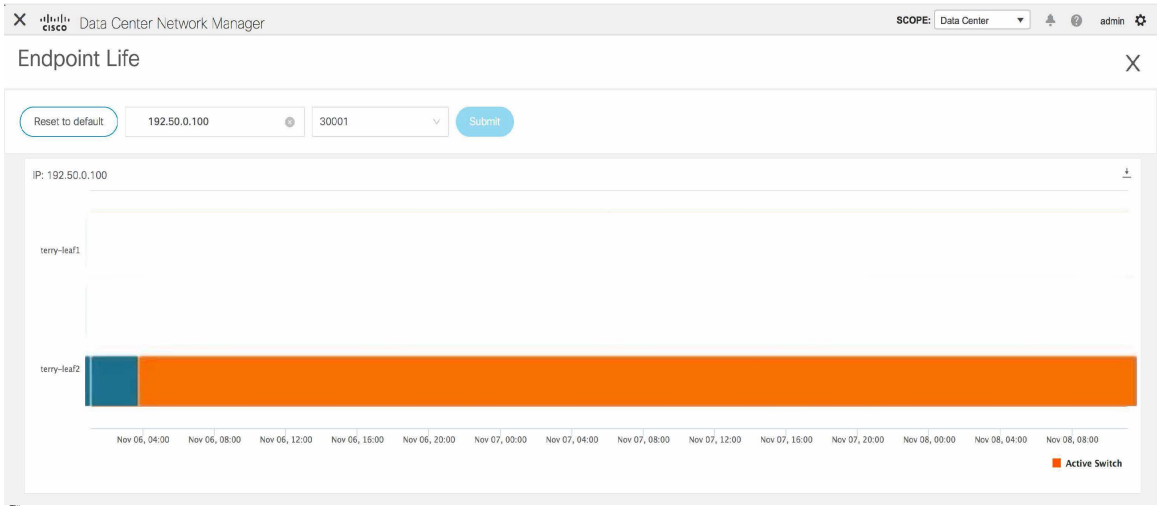

# Endpoint Life

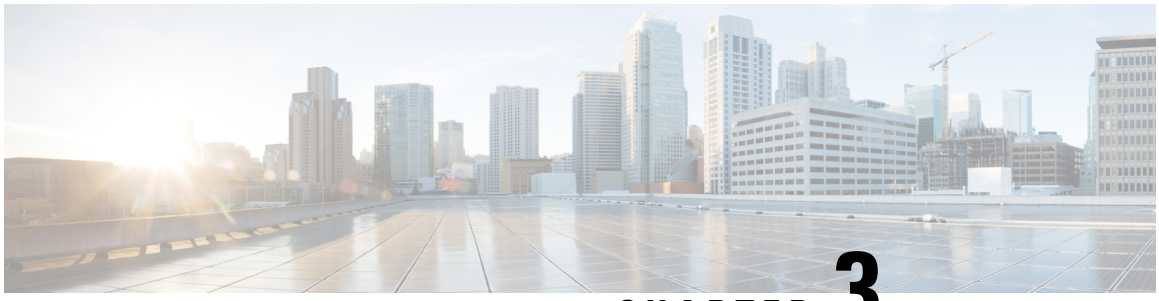
Click **Actions** > **Endpoint Life** to display a time line of a particular endpoint in its entire existence within the fabric.

Specify the IP or MAC address of an endpoint and the VXLAN Network Identifier (VNI) to display the list of switches that an endpoint was present under, including the associated start and end dates. Click **Submit**.

Initiate a search by using an IPv4 or IPv6 address to display the **Endpoint Life** graph for IPv4/IPv6 endpoints. Initiate a search by using a MAC address to display the **Endpoint Life** graph for MAC-Only endpoints.

The window that is displayed is essentially the endpoint life of a specific endpoint. The bar that is orange in color represents the active endpoint on that switch. If the endpoint is viewed as active by the network, it will have a band here. If an endpoint is dual-homed, then there will be two horizontal bands reporting the endpoint existence, one band for each switch (typically the vPC pair of switches). In case the endpoints are deleted or moved, you can also see the historical endpoint deletions and moves on this





CHAPTER 3

# Topology

**UI Navigation** - Click **Topology**.

The **Topology** window displays color-encoded nodes and links that correspond to various network elements, including switches, links, fabric extenders, port-channel configurations, virtual port-channels, and more. Use this window to perform the following tasks:

- To view more information about each of these elements, hover your cursor over the corresponding element.
- To view your navigation in the topology, view the breadcrumb at the top.
- When you click the device or the element, a slide-in pane appears from the right that displays more information about the device or the element. To view more information in the topology, double-click a node to open the node topology. For example, to view the fabric topology and its components in the **Topology** window, double-click the fabric node and then double-click an element that you want to view such as a host, a multicast group or a multicast flow, as applicable to the fabric type, and view the respective topology.
- If you want to view the fabric summary for the fabrics, click the fabric node. From the **Fabric Summary** slide-in pane, open the **Fabric Overview** window. Alternatively, you can right-click a fabric and choose **Detailed View** to open the **Fabric Overview** window. For more information about fabric overview window, see Fabric Overview, on page 181.
- Similarly, you can click on a switch to display the configured switch name, IP address, switch model, and other summary information such as status, serial number, health, last-polled CPU utilization, and last-polled memory utilization in the **Switch** slide-in pane. To view more information, click the **Launch** icon to open the **Switch Overview** window. For more information about switch overview window, see Switches, on page 279.
- Choose an action from the **Actions** drop-down list to perform various actions based on the element you select in the topology.

  For example, when you open the data center topology view, the only action available in the actions drop-down list is Add Fabric. However, when you open the fabric topology view, many more options are available in the drop-down list. For example, for LAN fabrics, the available actions are Detailed View, Edit Fabric, Add Switches, Recalculate Config, Preview Config, Deploy Config, Add Link, Deployment Disable, Backup Fabric, Restore Fabric, VXLAN OAM, and Delete Fabric.. Note that for IPFM fabrics, the available actions are Detailed View, Edit Fabric, Add Switches, Recalculate Config, Preview Config, Deploy Config, and Delete Fabric.

- To perform actions on the elements in the topology, other than the ones listed in the actions drop-down list, right-click the element. This opens the appropriate windows and allows you to perform tasks based on the elements. For example, if you right-click a fabric, you can perform tasks such as various configurations, delete the fabric, backup and restore, and many more.
- The VXLAN OAM option appears in the **Actions** drop-down list only for VXLAN Fabric, eBGP VXLAN Fabric, External, and Lan Classic fabric technologies, which support VXLAN OAM. For more instructions, see VXLAN OAM, on page 164.

The IPFM fabric topology is specific to the operations performed by Nexus Dashboard Fabric Controller IP for Media Fabric (IPFM) and applicable for both the IPFM and Generic Multicast modes .



**Note** In a flow topology that involves the Ingress and Egress nodes, the arrows in the node icon indicate the direction of the flow from the Ingress node or sender (indicated by **(S)**) to the Egress node or receiver (indicated by **(R)**).

This section contains the following:


- Searching Topology, on page 18
- Viewing Topology, on page 19


# Searching Topology

Use a combination of search attributes and search criteria in the search bar for an effective search. As you enter a combination of search attribute and search criteria in the search bar, the corresponding devices are highlighted in the topology.

You can apply the search criteria such as equals (**=**), does not equal (**!=**), contains (**contains**), and does not contain (**!contains**).

The search attributes that you can use for LAN fabrics are ASN, Fabric Type, Fabric Name, and Fabric technology. The fabric type attributes that you can use for search include switch fabric, multi-fabric domain, external, and LAN monitor. The fabric technology attributes that you can use for search include fabricpath fabric, VXLAN fabric, VLAN fabric, external, LAN classic, IPFM classic, IPFM fabric, switch group, multi-fabric domain, eBGP VXLAN fabric, eBGP routed fabric, MSO site group, meta fabric, LAN monitor fabric, and IOS-XE VXLAN fabric.

For IPFM fabrics, the following fields are available to search on: switch or hostname, switch or host IP address, switch MAC, and switch serial number. In the Generic Multicast mode, also, you can search the receiver-interface name or IP addresses in this window.

When a device is displayed on the topology, double-click it to navigate further into the topology. For example, when the fabric that you searched is displayed on the topology, double-click on the fabric (cloud icon) to navigate inside its topology. Furthermore, after the fabric is displayed on the topology, you can continue to search based on a combination of a criteria and various search attributes such as VPC peer, IP address, model, mode, switch, switch role, discovery status, software version, up time, and serial.

**Note** Certain levels of the topology allow filters only, that is, filters take the place of Search. The topology listing for these levels display a limited number of entities. For example, Easy Fabric Networks are limited to 50 networks shown. Filters must be used to see additional elements or entities.

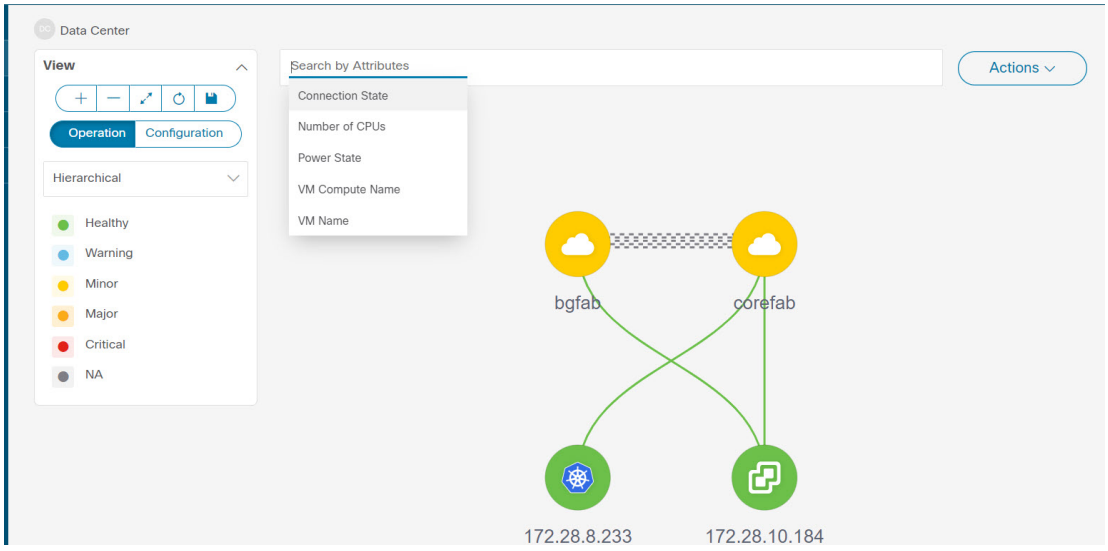
# Viewing Topology

To pan, click and hold anywhere in the whitespace and drag the cursor up, down, left, or right. To drag switches, click, hold, and move the cursor around the whitespace region of the topology.

In case of multiple selection of switch, you must release the modifier keys (cmd/ctrl) before releasing mouse drag to end the switch selection.

You can view the following information of the devices and links in the **View** pane:

- Layout options - You can zoom in, zoom out, or adjust the layout to fit the screen. You can also refresh the topology or save any changes to the topology. For more information, see Zooming, Panning, and Dragging, on page 34.
- Logical Links - For LAN topologies, you can view the logical links using the **Show Logical Links** toggle switch.
- Operation/Configuration - For LAN topologies, you can also select operation or configuration.
- Select Layout drop-down list - Choose the layout for your topology from this drop-down list, and click **Save Topology Layout** in the layout options. For more information, see Layouts, on page 34.
- Status - The status of every device or link is represented by different colors. You can view the configurational status and operational status as well for LAN topologies. For more information, see Status, on page 34.

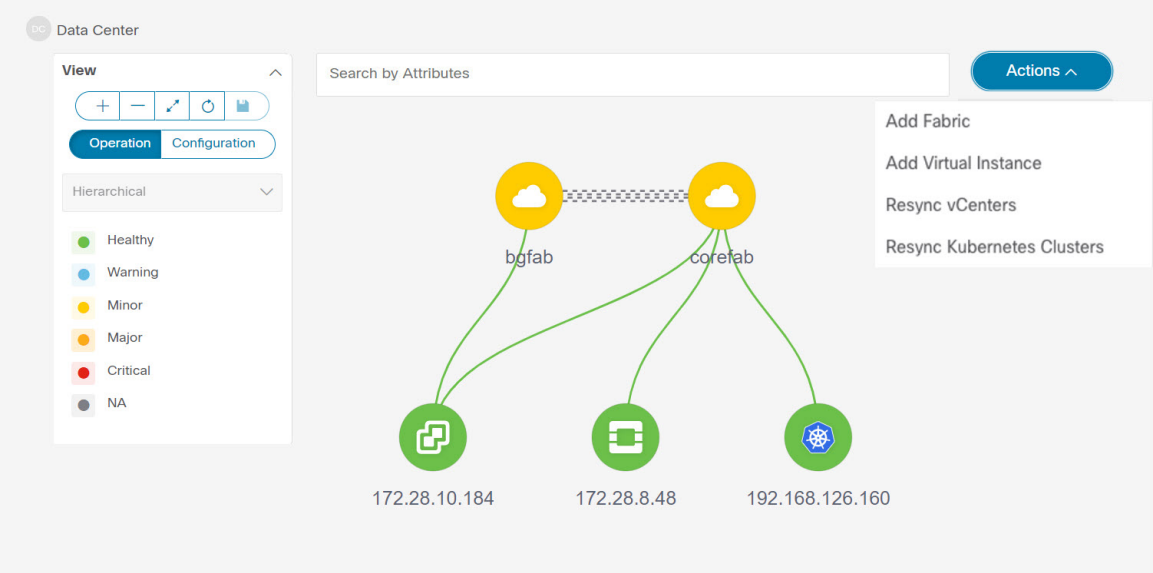



Topology for a node is displayed at multiple scope. Each scope is shown in the hierarchical order. The scope hierarchy is shown as breadcrumbs and can be navigated to required scope. Scopes are as follows:

- Data Center
- Cluster (vCenter)
- Resource List (DVS, Compute, and VM)
- Resource

**Note**

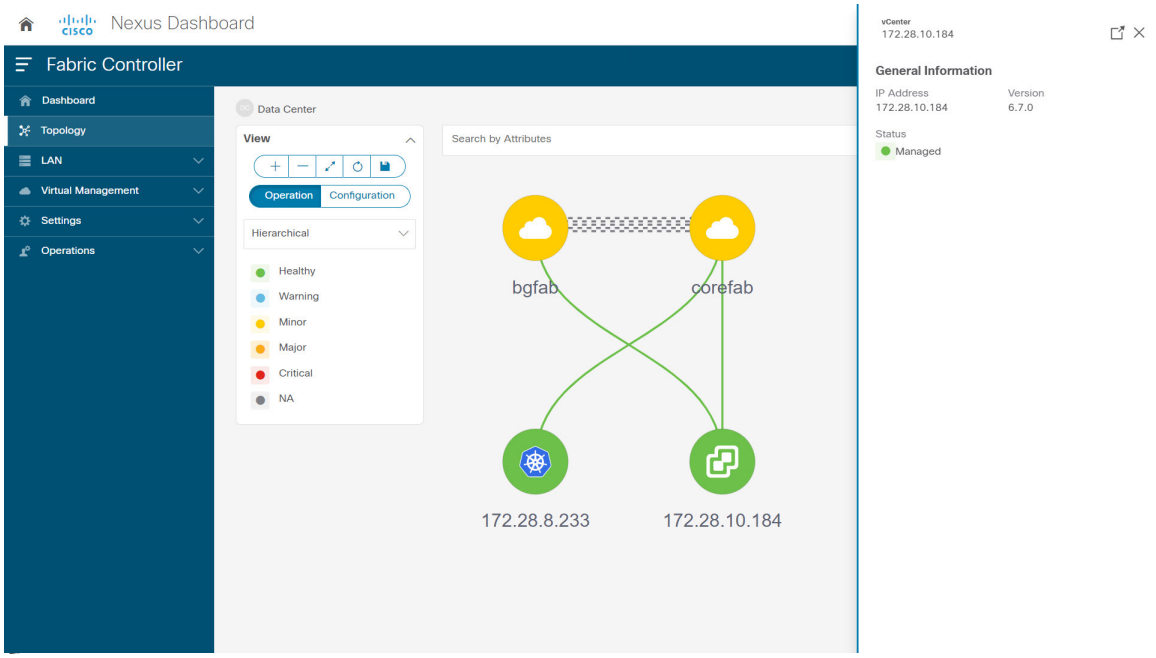
- In the **Topology** window, FEX appears in gray (**Unknown** or **NA**) because Operation and Configuration status is not calculated for FEX.
- After moving a cable from one port to another port, the old fabric link is retained in the **Topology** window, and it is shown in the red color indicating that the link is down. Right-click on the link and delete it if the removal was intentional. A manual Rediscover of the switch will also delete and re-learn all links to that switch.



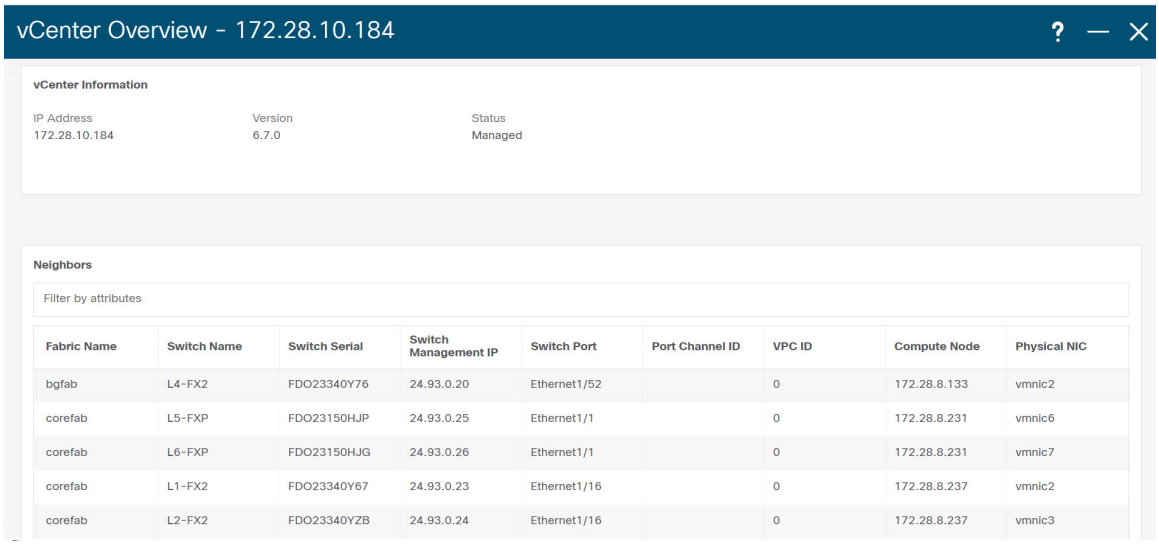

# Viewing vCenter Visualization

In a virtualized environment, troubleshooting is intiated with identifying network attachment point for Virtual Machines (VMs). This process discovers critical details such as server, virtual switch, port group, VLAN, associated network switch, and physical port. These requires multiple touch points and communication between server, network administrator and other applications like compute orchestrator, compute manager, network manager, network controller.

Click on the vCenter visualization node, a slide-in panel appears, click on **Launch** icon to view vCenter Overview



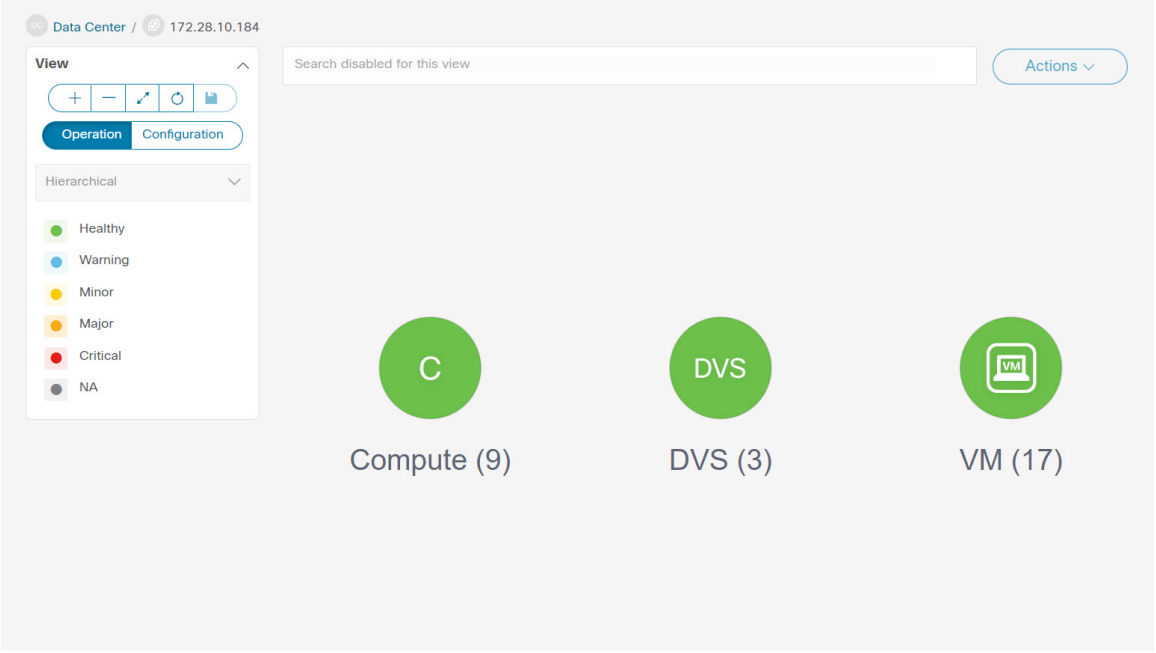

This window has summarized data such as vCenter IP address, status of vCenter, fabric associated with the cluster, Switch name, Switch IP, Switch Port, VPC ID, Compute Node and Physical



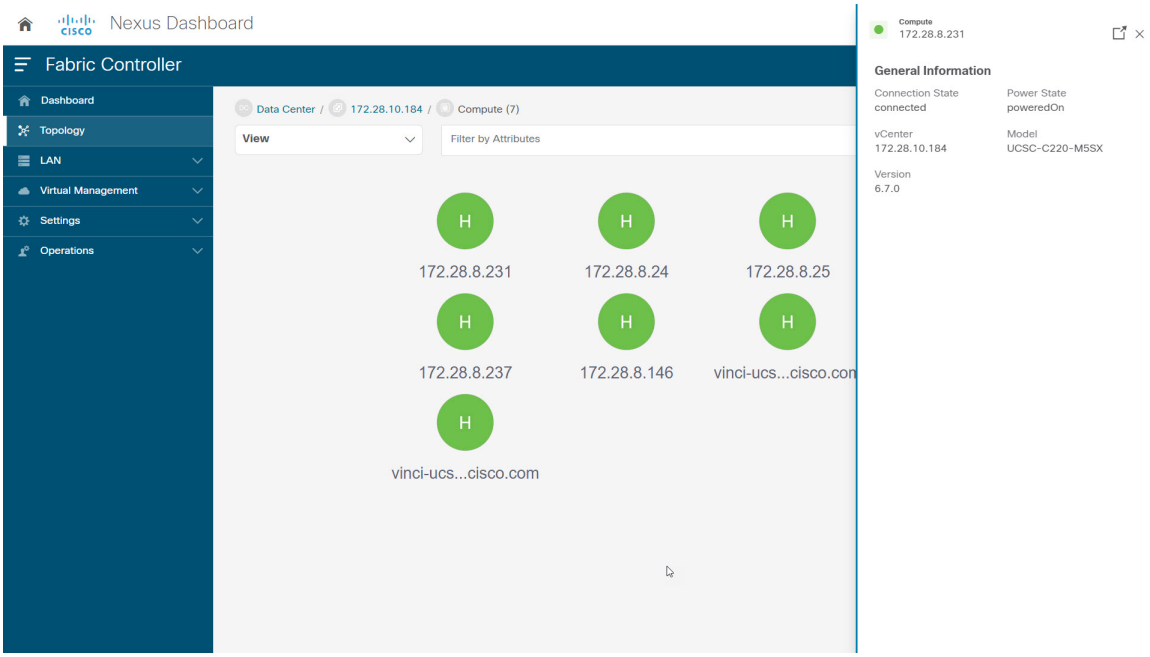

Double-click on the vCenter cluster node to view the associated vCenter cluster resources such as Compute, DVS, VMs. Each node has a number displayed in brackets, which indicates the number of specific nodes in the vCenter instance.

Double-click on Compute, or DVS, or VMs to view required list of resource type and its topology.

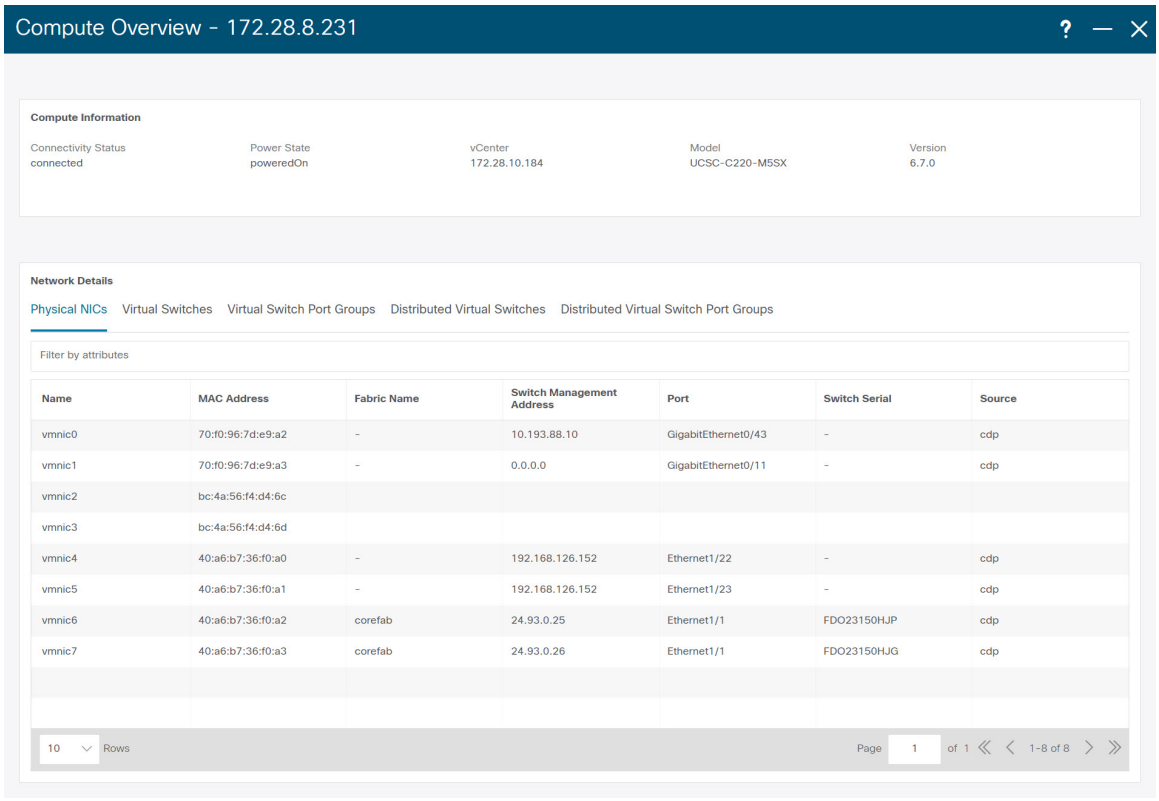
**Note** When you double-click on DVS, it displays the associated compute hosts under the DVS.



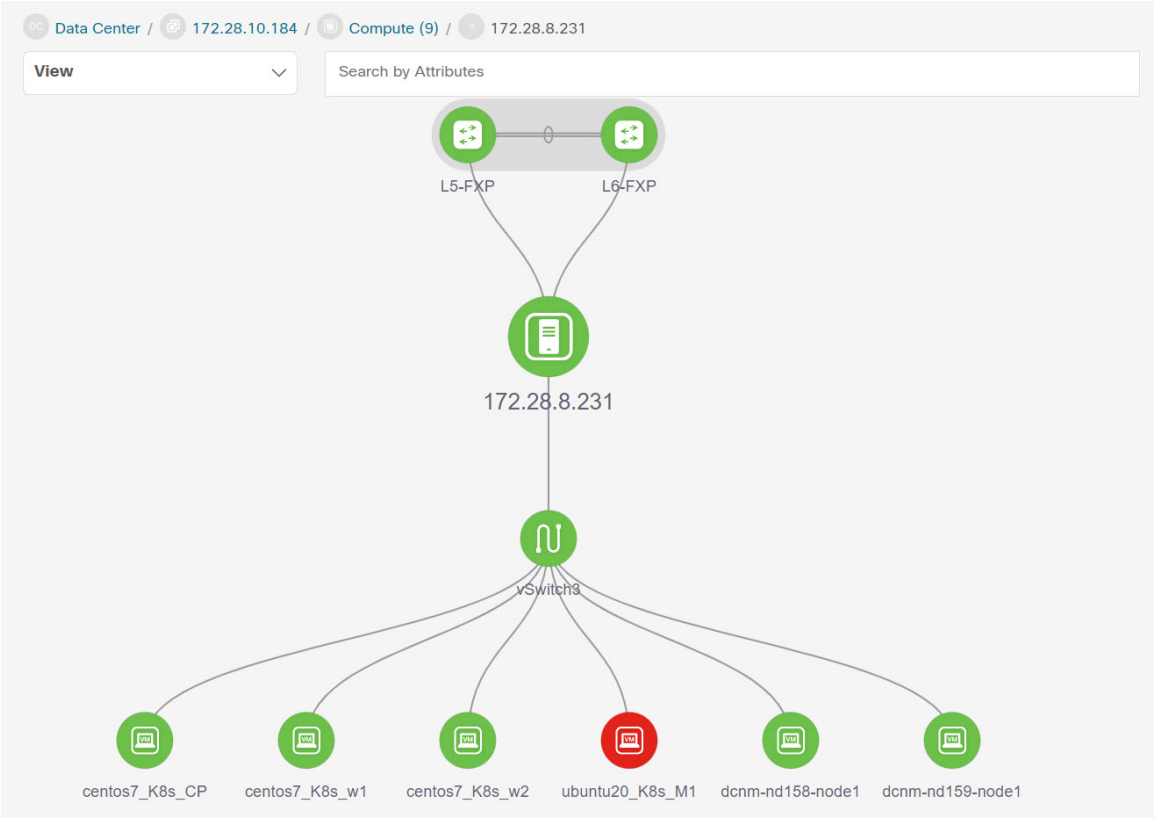

Click on a node, a slide-in panel appears, click on **Launch** icon to view **Compute Overview** window.



You can view the Compute information and Network details tabs which displays information such as power state, memory size, IP address, MAC address associated with the node.



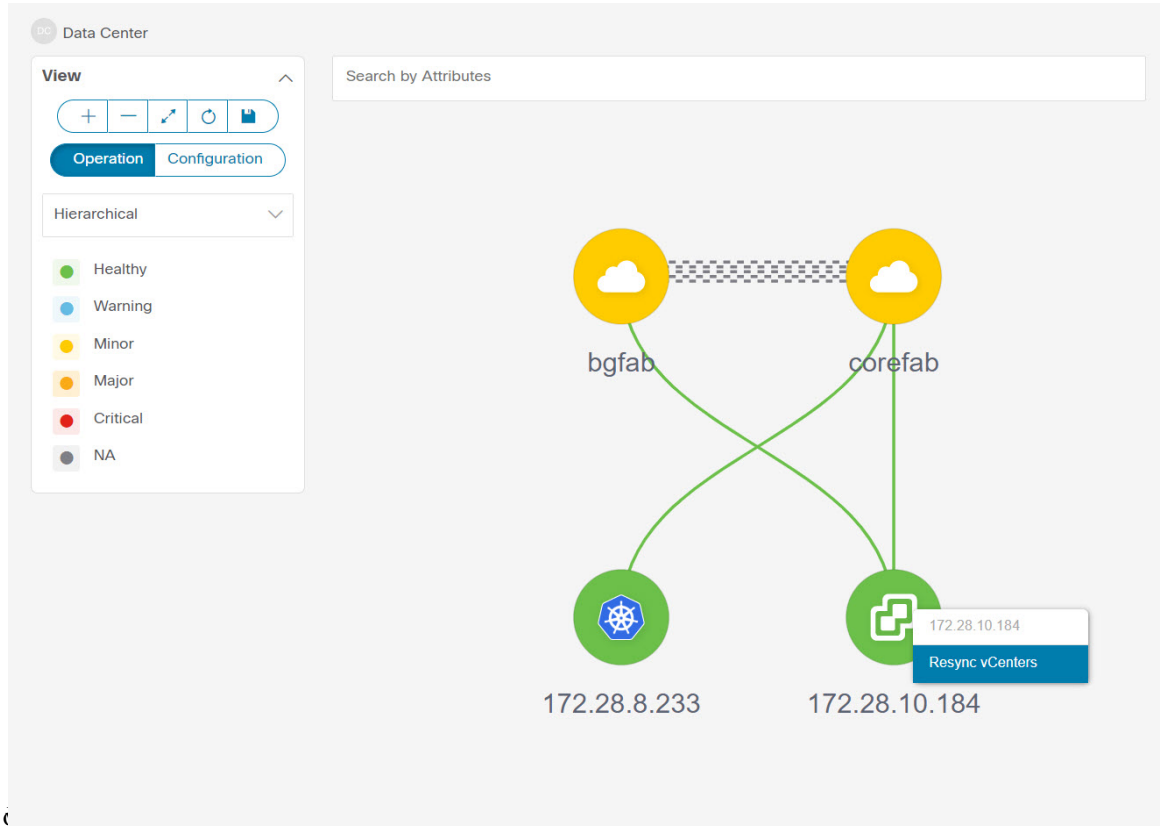
You can search using **Search by Attributes** to search required node. Double-click on the specific node to view the complete topology of vCenter




## Resync vCenter

Resync synchronizes the state of all on board vCenter clusters. To resynchronize vCenter clusters, right-click on topology window, choose **Resync vCenters** and click **Confirm**. To synchronize individual vCenter cluster,

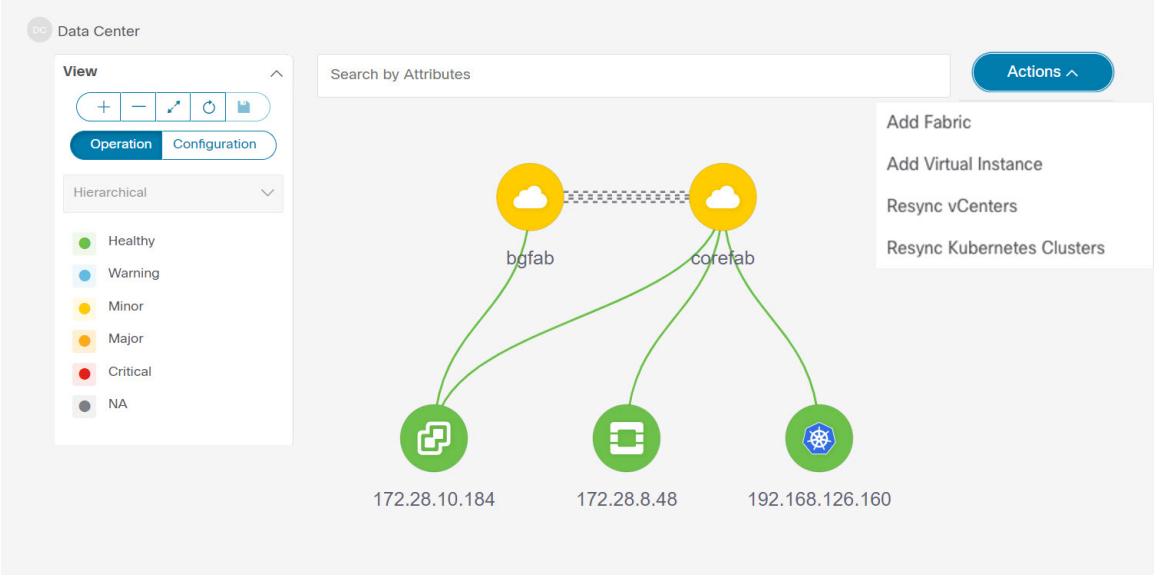
choose the Rediscover




# Viewing Kubernetes Cluster

You can view topology in multiple scope, each scope is displayed in the hierarchical order and navigation breadcrumb. These scopes are:

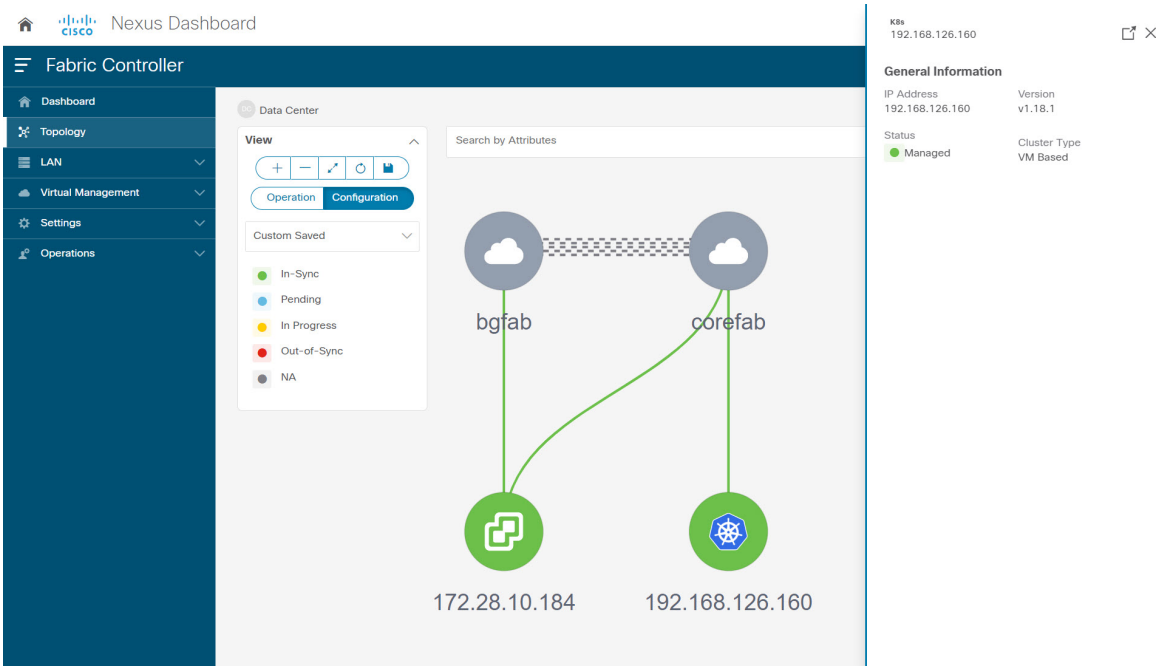
- Data Center, Cluster (Kubernetes)
- Resource List (Compute, and Pod)
- Resource (Compute and Pod)



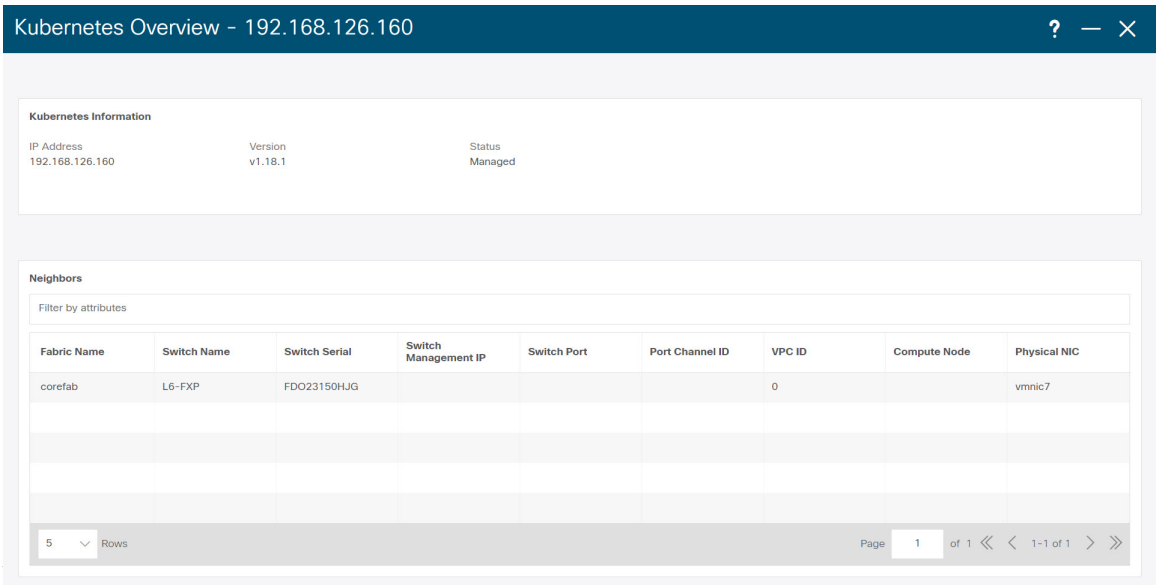

Kubernetes Clusters are of two types:

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

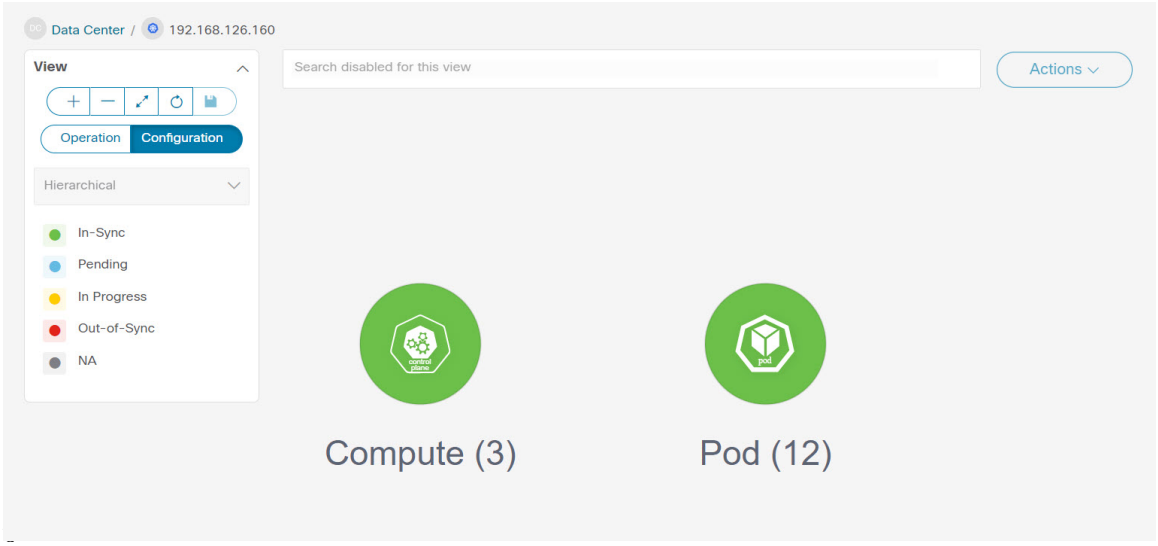
Click on the Kubernetes cluster node, a slide-in panel appears, click on **Launch** icon to view **Kubernetes Overview**



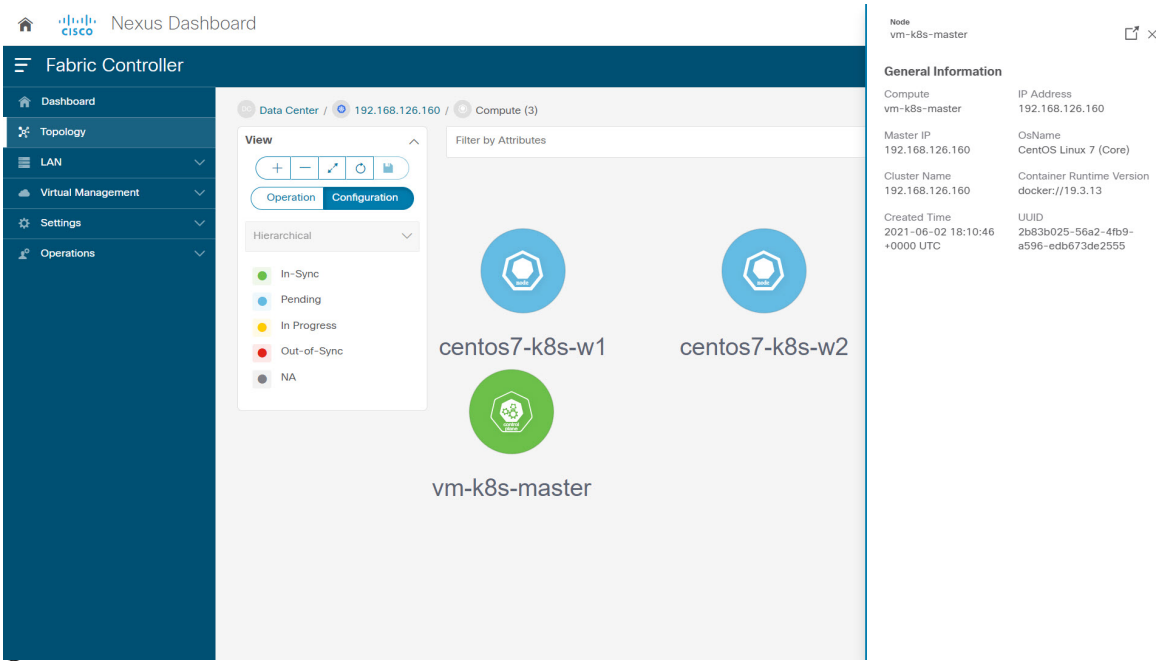
This window has summarized data such as vCenter IP address, status of vCenter, fabric associated with the cluster, Switch name, Switch IP, Switch Port, VPC ID, Compute Node and Physical



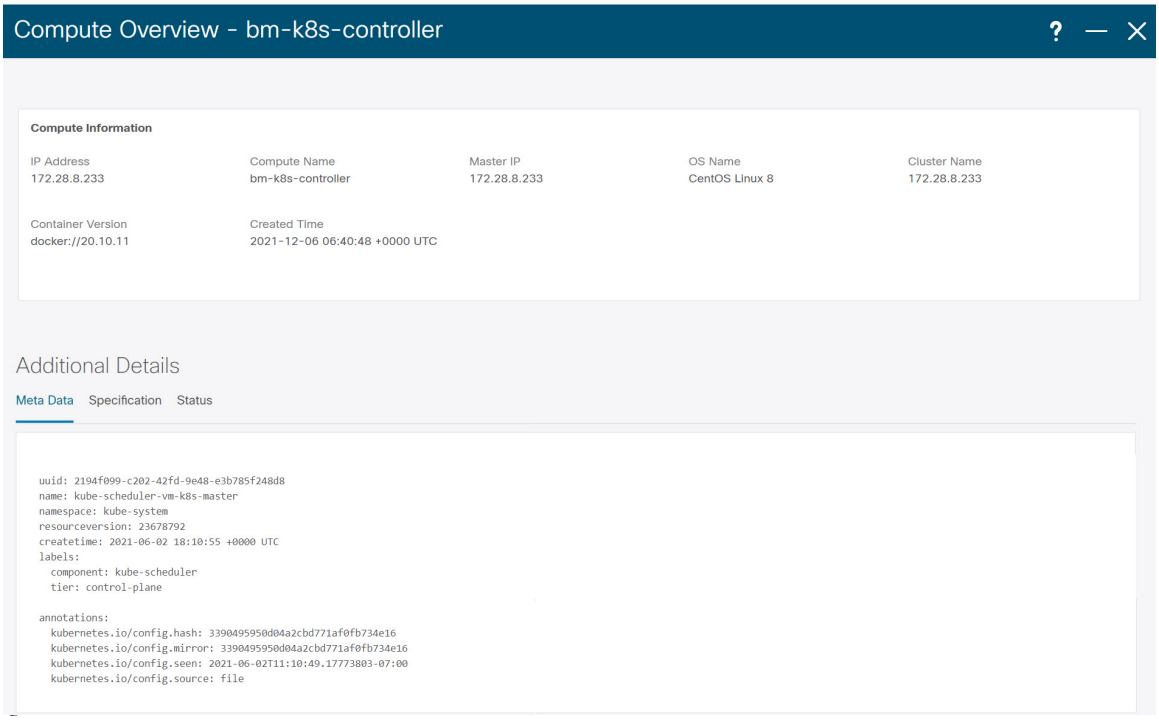

Double-click on the Kubernetes cluster node to view the associated Kubernetes cluster resources such as Computes and Pods. Each node as a number displayed in brackets, which indicates the number of specific nodes in the Kubernetes



Double-click on appropriate resource (computes or pods) group to display the list of computes and the pods in the Kubernetes cluster. You can search the specific node using **Filter by**



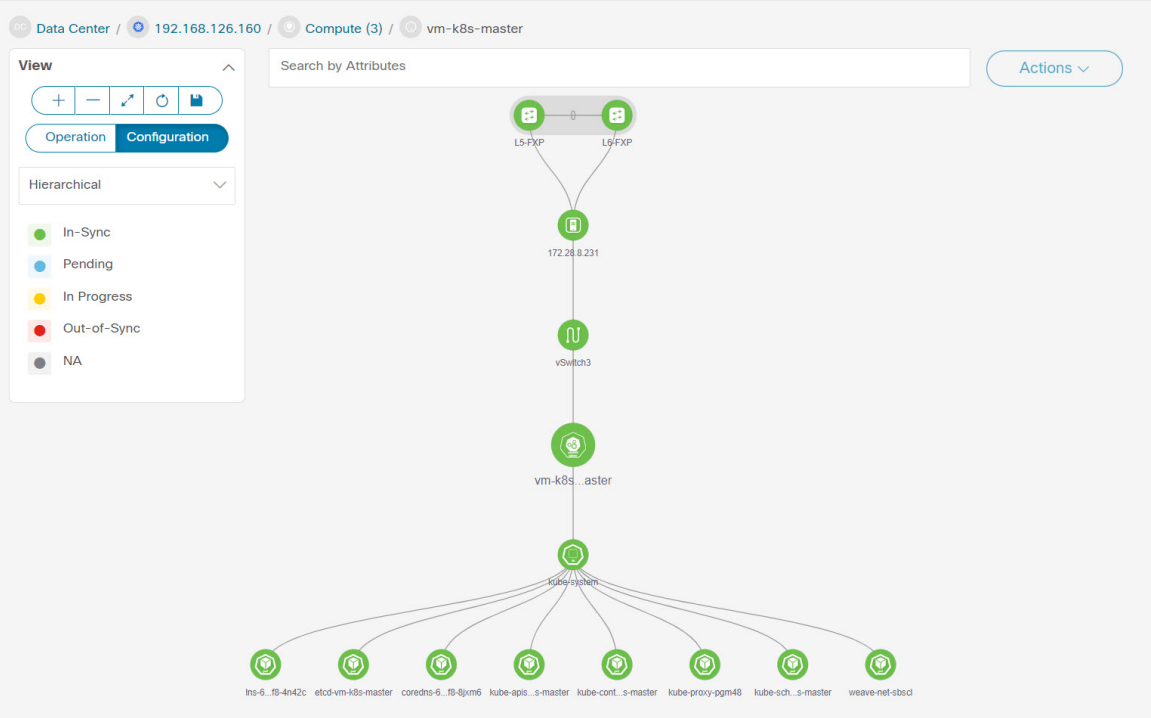

Click on the Nodes to view details about the node. A side panel appears, showing the Node Summary. Click **Launch** icon to view Meta Data, Specifications, and Status information for the selected




Meta data tab consists of Kubernetes node or Pod name. Specification tabs include the desired design or configuration of the node or the Pod. Status tab indicates the running state information of the node or the pod.

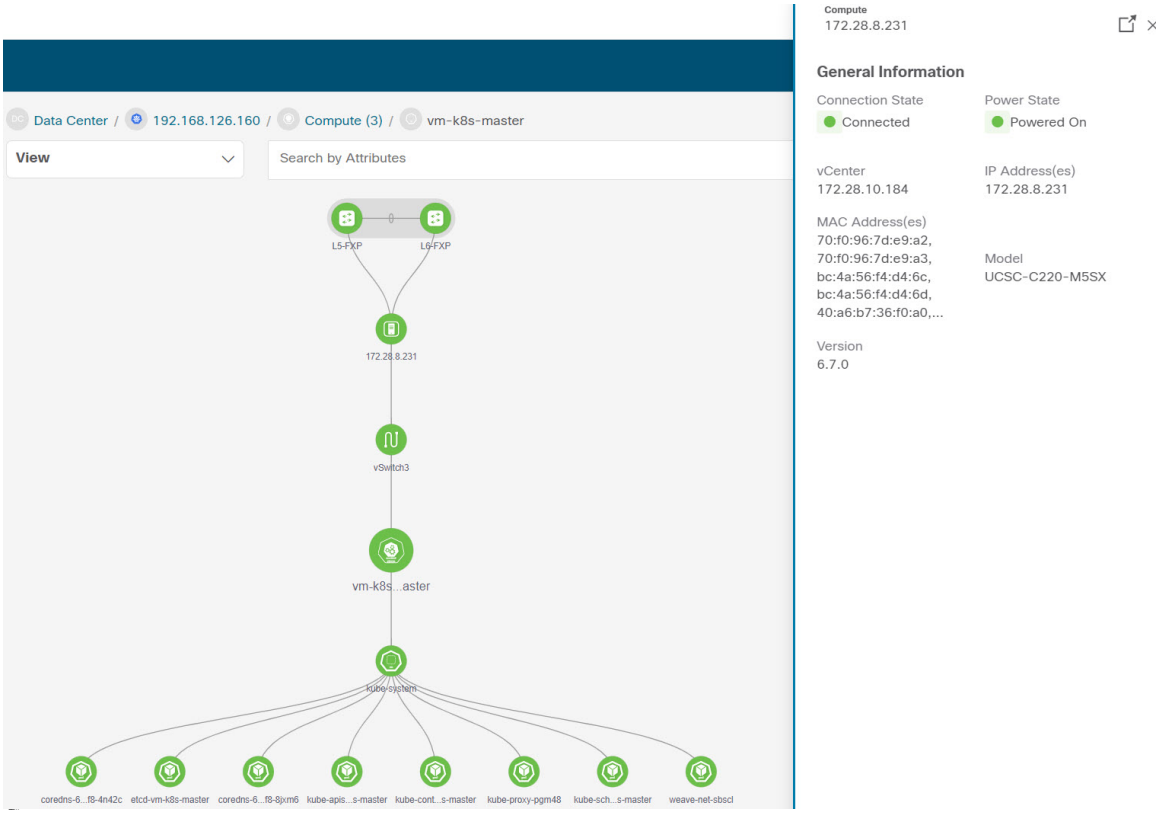
Click on Compute or Pod to view specific compute or pod node details. You can search using **Filter by Attributes** to search required node.

Double-click on the specific node to view the complete topology of vCenter



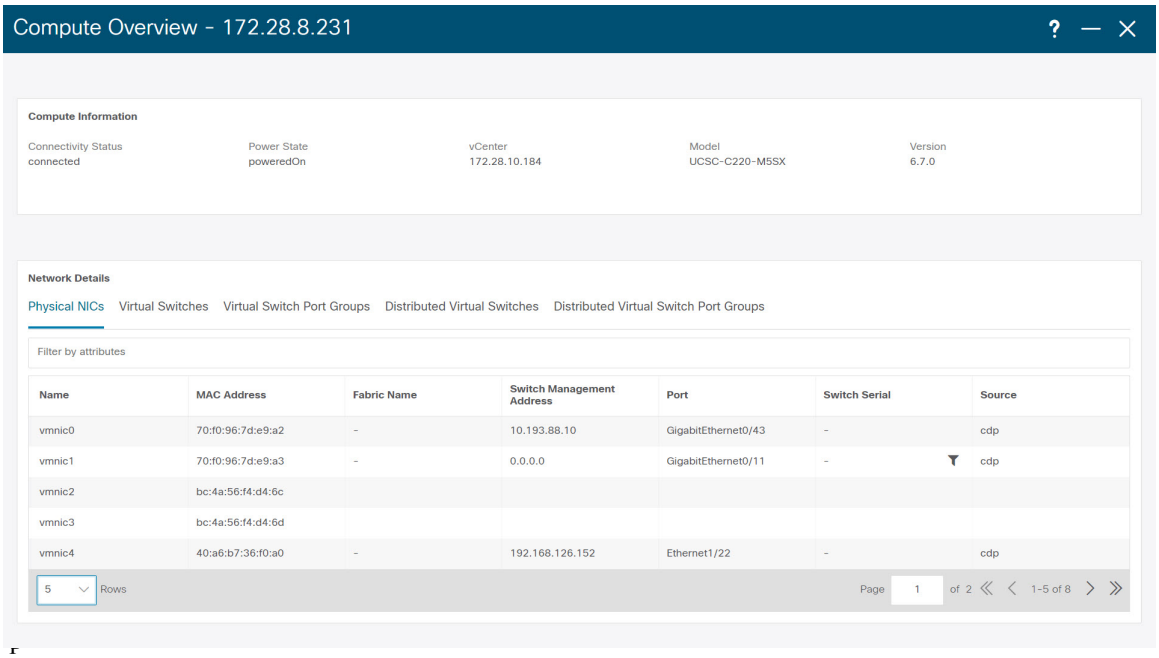

Click on a cluster node, a slide-in panel appears, click on **Launch** icon to view Kubernetes Cluster Node Overview window. To view the Compute information and Network details tabs.

Click on the pod node, a slide-in panel appears. Click on **Launch** icon to view the Kubernetes Pod Overview




**Compute Information** – Displays connectivity status, Power state, vCenter IP, Model and Version.

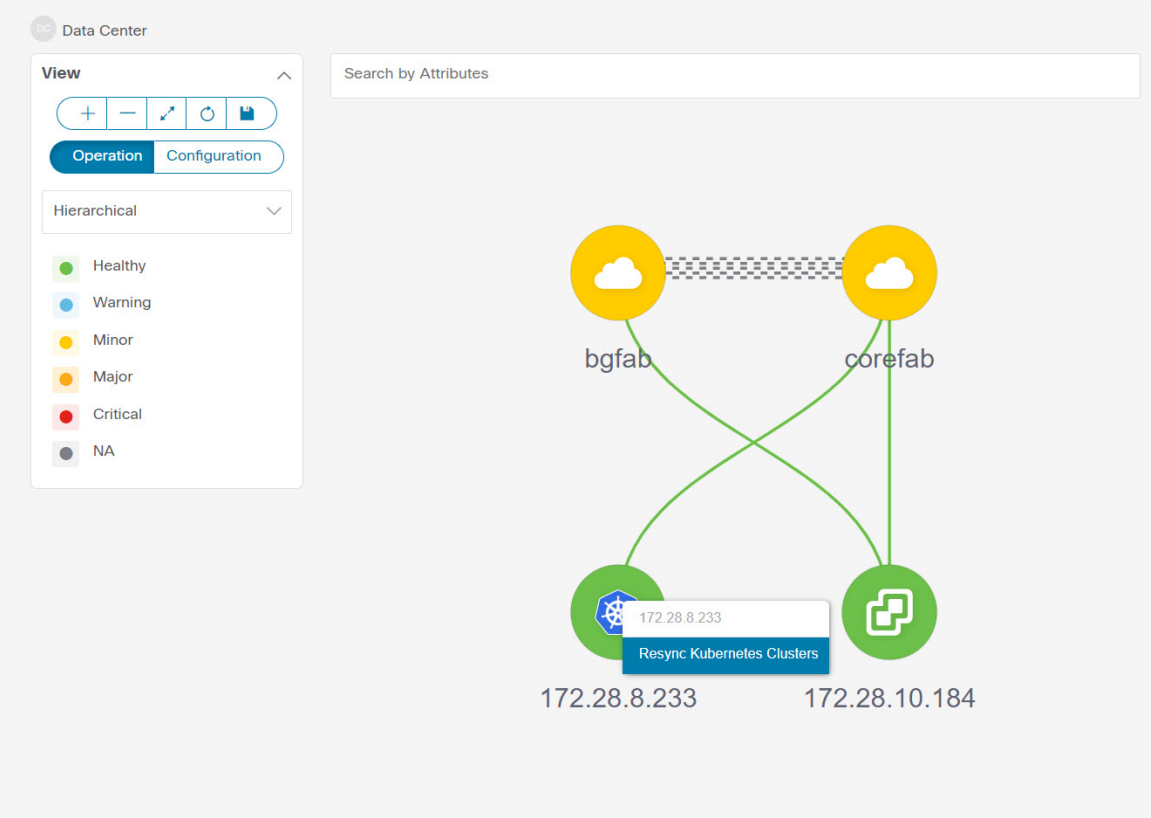
**Network Details** – Displays tabular information such as Physical NICs, Virtual Switches, Virtual Switch Port Groups, Distributed Virtual Switches, Distributed Virtual Switch Port



## Resync Kubernetes Clusters

To resynchronize kubernetes clusters, right-click on topology window, click **Resync Kubernetes Clusters** and click **Confirm**.

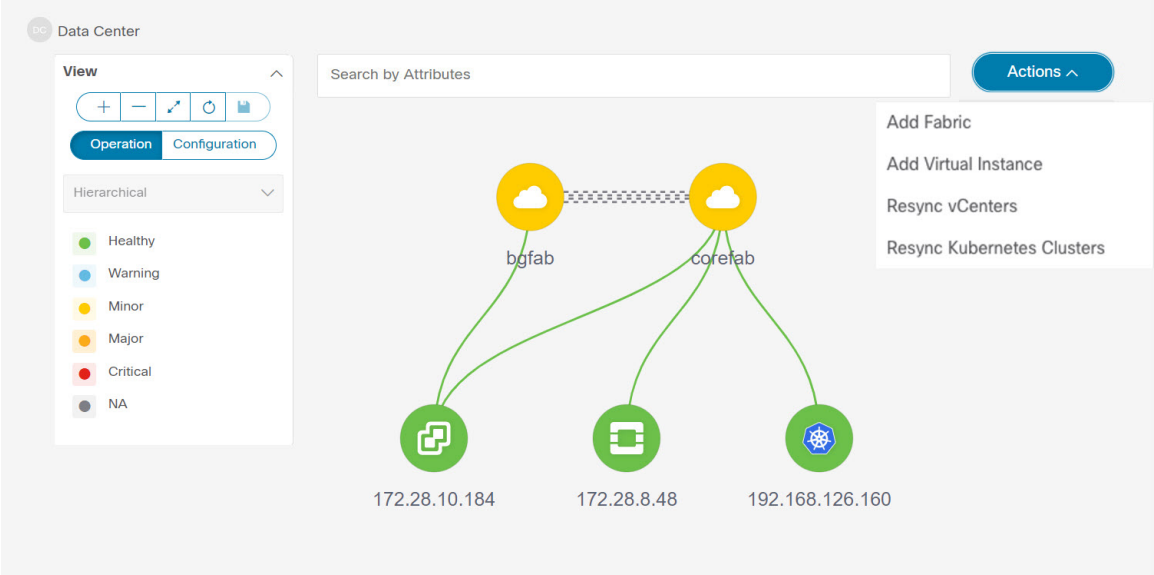
Resync synchronizes the state of all onboarded Kubernetes clusters.




# Viewing OpenStack Cluster

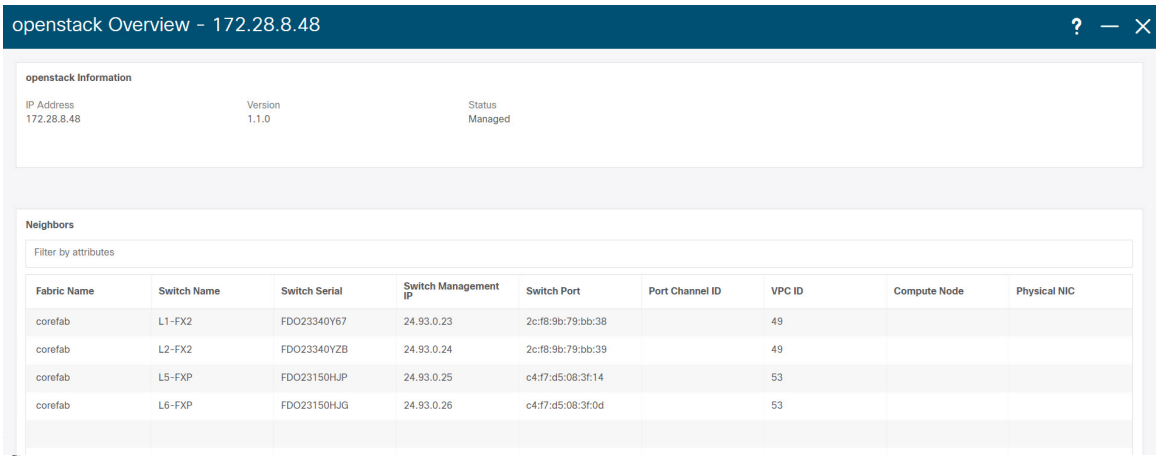
Topology for a node is displayed at multiple scope. Each scope is shown in the hierarchical order. The scope hierarchy is shown as breadcrumbs and can be navigated to required scope. Scopes are as follows:

- Data Center
- Cluster (Openstack)
- Resource List (Compute, and VM)
- Cluster

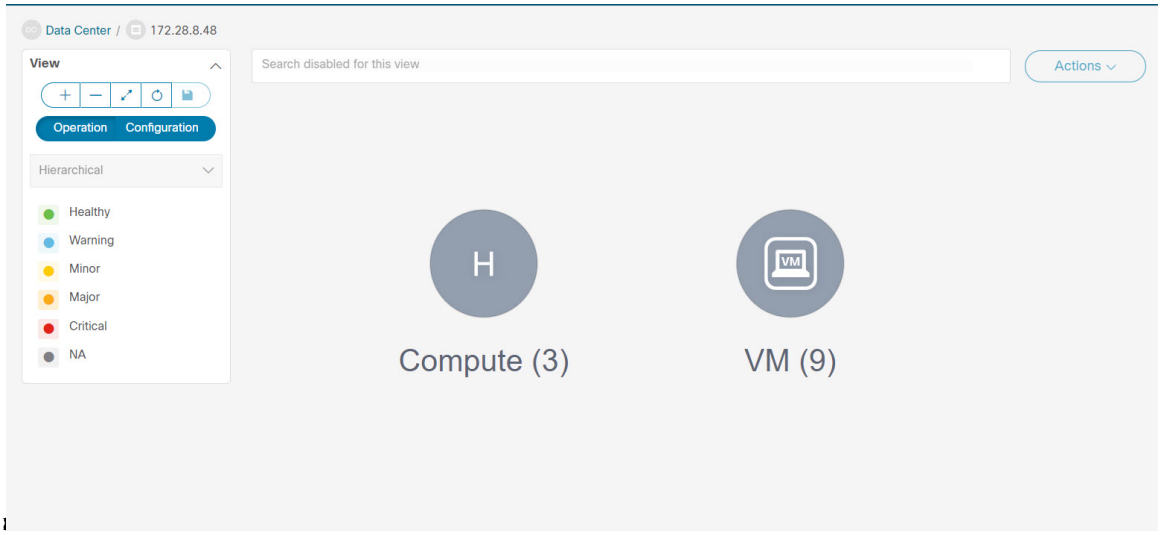



Click on the Openstack cluster node, a slide-in panel appears, click on **Launch** icon to view Openstack cluster window.

This window has summarized data such as Openstack cluster IP address, status of vCenter, fabric associated with the cluster, Switch name, Switch IP, Switch Port, VPC ID, Compute Node and Physical



Double-click Openstack cluster node, to view associated VMs and compute nodes. Each node has a number displayed in brackets, which indicates the number of specific nodes in the vCenter




Double-click on Compute or VM group icon to view list of specific compute or VMs in the cluster.

You can search using **Filter by Attributes** to search required node.

Double-click on the specific node to view the complete topology of Openstack cluster node.

# IPFM - Multicast Flow

Generic Multicast is not limited to the two-tier spine or leaf topology. The flow classification and path tracing are not limited to any specific topology if all the involved switches are Cisco Nexus 9000 Series switches with the Cisco NX-OS Release 9.3(5). Generic Multicast is supported for the default VRF.

**Note**

- If you remove a device from the Inventory, the Policy deployment status for that switch is removed. However, clear the policy configuration on the switch also.

To enable prefix for multicast, perform the following steps:

1. From Nexus Dashboard Fabric Controller Web UI, choose **Settings** > **Server Settings**.
2. Click **IPFM** tab, and check the check box **Enable mask/prefix for the multicast range in Host Policy**.
3. Click **Save**.

To view the multicast flows topology, perform the following steps:

1. Double-click the IPFM fabric in the **Topology** window.
2. Double-click the Multicast Flows node.
3. Double-click the required Multicast Flow.

   The multicast flow topology is displayed.

A multicast flow topology involves spine, leaf, and sender and receiver hosts. The dotted moving lines depict the flow of traffic in the IPFM fabric topology. The arrows in the icon indicate the direction of the flow, and the IP address suffixed with **(S)** and **(R)** indicate the sender and receiver host respectively.

# Zooming, Panning, and Dragging

You can zoom in and zoom out using the controls that are provided at the bottom left of the windows or by using your mouse's wheel.

To pan, click and hold anywhere in the whitespace and drag the cursor up, down, left, or right.

To drag switches, click, hold, and move the cursor around the whitespace region of the topology.

# Layouts

The topology supports different layouts along with a **Save Layout** option that remembers how you positioned your topology.

- **Hierarchical** and **Hierarchical Left-Right** - Provide an architectural view of your topology. Various switch roles can be defined that will draw the nodes on how you configure your CLOS topology.



**Note** When running a large-scale setup, being able to easily view all your switches on a leaf-tier can become difficult. To mitigate this, Nexus Dashboard Fabric Controller splits your leaf-tier every 16 switches.

- **Circular** and **Tiered-Circular** - Draw nodes in a circular or concentric circular pattern.
- **Random** - Nodes are placed randomly on the window. Nexus Dashboard Fabric Controller tries to make a guess and intelligently place nodes that belong together in close proximity.
- **Custom saved layout** - Nodes can be dragged around according to your preference. After you position as required, click **Save** to retain the positions. The next time you come to the topology, Nexus Dashboard Fabric Controller will draw the nodes based on your last saved layout positions.

Before a layout is chosen, Nexus Dashboard Fabric Controller checks if a custom layout is applied. If a custom layout is applied, Nexus Dashboard Fabric Controller uses it. If a custom layout is not applied, Nexus Dashboard Fabric Controller checks if switches exist at different tiers, and chooses the Hierarchical layout or the Hierarchical Left-Right layout. Force-directed layout is chosen if all the other layouts fail.

# Status

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

- Green - Indicates that the element is in good health and functioning as intended.
- Blue - Indicates that the element is in a warning state and requires attention to prevent any further problems.
- Yellow - Indicates that the element has minor issues.

- Orange - Indicates that the element has major issues and requires attention to prevent any further problems.
- Red - Indicates that the element is in critical state and requires immediate attention.
- Gray: Indicates lack of information to identify the element or the element has been discovered.

The configurational colors and what they indicate are described in the following list:

- Green - Indicates that the element is element is In-Sync with the intended configuration.
- Blue - Indicates that the element has pending deployments.
- Yellow - Indicates that active deployments are in-progress.
- Red - Indicates that the element is Out-of-Sync with the intended configuration.
- Gray: Indicates lack of information or no support for Configuration Sync calculation.



**Note**

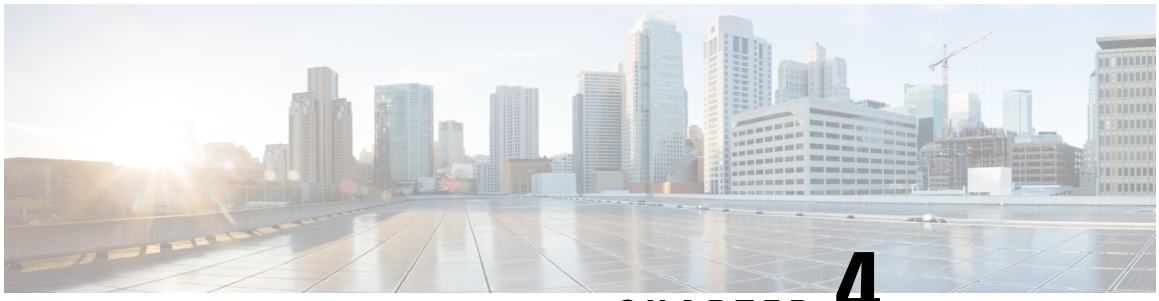
- In the **Topology** window, FEX appears in gray (**Unknown** or **n/a**) because Operation and Configuration status is not calculated for FEX.
- After moving a cable from one port to another port, the old fabric link is retained in the **Topology** window, and it is shown in the red color indicating that the link is down. Right-click on the link and delete it if the removal was intentional. A manual Rediscover of the switch will also delete and re-learn all links to that switch.



PART I

# LAN


- Fabrics, on page 39
- Switches, on page 279
- Policies, on page 313
- Interfaces, on page 317



CHAPTER 4

# Fabrics


- LAN Fabrics, on page 39
- Enhanced Role-based Access Control, on page 154
- Enhanced RBAC Use-Cases, on page 158
- Nexus Dashboard Security Domains, on page 161
- Backup Fabric, on page 163
- Restoring Fabric, on page 163
- VXLAN OAM, on page 164
- Endpoint Locator , on page 166
- Fabric Overview, on page 181


## LAN Fabrics

The following terms are referred to in this 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 Nexus Dashboard Fabric Controller using the Data Center VXLAN EVPN fabric template.
  - NFM migration to Cisco Nexus Dashboard Fabric Controller using the Data Center VXLAN EVPN fabric template.

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

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

The following table describes the fields that appear on **LAN** > **Fabrics**.

| Field | Description |
|---|---|
| Fabric Name | Displays the name of the fabric. |
| Fabric Technology | Displays the fabric technology based on the fabric template. |

| Field | Description |
|---|---|
| Fabric Type | Displays the type of the fabric—Switch Fabric, LAN Monitor, or External |
| ASN | Displays the ASN for the fabric. |
| Fabric Health | Displays the health of the fabric. |

The following table describes the action items in the Actions menu drop-down list, that appear on **LAN** > **Fabrics**.

| Action Item | Description |
|---|---|
| Create Fabric | From the Actions drop-down list, select **Create Fabric**. For more instructions, see Create a Fabric, on page 45. |
| Edit Fabric | Select a fabric to edit. From the Actions drop-down list, select **Edit Fabric**. Make the necessary changes and click **Save**. Click **Close** to discard the changes. |
| Delete Fabric | Select a fabric to delete. From the drop-down list, select **Delete Fabric**. Click **Confirm** to delete the fabric. |

# Fabric Summary

Click on a fabric to open the side kick panel. The following sections display the summary of the fabric:

- **Health** - Shows the health of the Fabric.
- **Alarms** - Displays the alarms based on the categories.
- **Fabric Info** - Provides basic about the Fabric.
- **Inventory** - Provides information about Switch Configuration and Switch Health.

Click the **Launch** icon to the right top corner to view the Fabric Overview.

# Understanding Fabric Templates

### Fabric Templates

The following table provides information about the available fabric templates:

| Fabric Template | Description | Detailed Procedures |
|---|---|---|
| Easy_Fabric | Fabric template for a VXLAN BGP EVPN deployment with an IGP (OSPF, IS-IS) and iBGP deployment for Nexus 9000 and Nexus 3000 switches. Both IPv4 and IPv6 underlay are supported. | Creating a VXLAN EVPN Fabric Using the Easy_Fabr Template, on page 49 |
| Easy_Fabric_IOS_XE | Fabric template for a VXLAN BGP EVPN deployment with Catalyst 9000 switches. | Create Easy_Fabric for Cisco Catalyst 9000 Series Switches, on page 101 |

| Fabric Template | Description | Detailed Procedures |
|---|---|---|
| Easy_Fabric_eBGP | Fabric template for an eBGP based Routed Fabric deployment with Nexus 9000 and Nexus 3000 switches. This template also supports an eBGP VXLAN BGP EVPN deployment with eBGP used as both the underlay and overlay protocol. | Creating VXLAN EVPN Fabric with eBGP-based Underlay, on page 539 |
| External_Fabric | Fabric template that supports Nexus and non-Nexus devices. Non-Nexus device support includes other Cisco devices (IOS-XE, IOS-XR) and third-party switches. This template can manage BGP configuration on core and edge routers.<br>Use cases for this template include:<br>• Classic hierarchical 2/3 Tier vPC or FabricPath like networks<br>• VXLAN EVPN deployments for Cisco Nexus switches other than Nexus 9000 or Nexus 3000<br>• VRF-Lite on Core/Edge devices<br>• Using NDFC in monitored mode, which is useful when you want to eventually move to managed mode after trying the monitor mode | Creating an External Fabric, on page 112 |
| LAN_Classic | Fabric template that allows you to monitor and manage various Cisco Nexus Greenfield and Brownfield deployments, including the traditional 2 or 3-tier, vPC, and FabricPath data center topologies. | LAN Fabrics, on page 39 |
| Fabric_Group | Fabric template that contains other LAN_Classic fabrics that allows visualization of groups of Classic LAN fabrics and their interconnections. | LAN Fabrics, on page 39 |
| LAN_Monitor | Fabric template to support Fabric Discovery persona for monitoring purposes only. Cisco Nexus Dashboard Insights (NDI) can operate on such a fabric. No configuration provisioning or image management is supported. | LAN Fabrics, on page 39 |
| MSD_Fabric | Fabric template for a VXLAN BGP EVPN Multi-Site Domain (MSD) deployment that can contain other VXLAN BGP EVPN fabrics with Layer-2/Layer-3 Overlay DCI Extensions. | Creating the VXLAN EVPN Multi-Site and Associ Member Fabrics, on page 593 |
| IPFM_Classic | Fabric template for existing deployments of IP Fabric for Media (IPFM). IPFM enables media content providers and broadcasters to use a flexible and scalable IP-based infrastructure. | Creating an IPFM Classic Fabric, on page 133 |
| Easy_Fabric_IPFM | Fabric template for an easy greenfield deployment of an IP Fabric for Media (IPFM) fabric. IPFM enables media content providers and broadcasters to use a flexible and scalable IP-based infrastructure. | Creating an IPFM Easy Fabric, on page 136 |

**Note** Enhanced Classic LAN is a preview feature in Nexus Dashboard Fabric Controller, Release 12.1.2e. We recommend that you use this feature marked as BETA in your lab setup only. Do not use this features in your production deployment.

To view Enhanced Classic LAN fabrics, you must enable this feature. On Web UI, navigate to **Settings** > **Server Settings** > **LAN-Fabric**, then check the **Enable Preview Features** check box.

| Type of Fabric | Description | REST API Template Name | Detailed Procedures |
|---|---|---|---|
| Data Center VXLAN EVPN | Fabric for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches. | Easy_Fabric | Creating a VXLAN EVPN Fabric Usin the Easy_Fabric Template, on page 49 |
| Enhanced Classic LAN | Fabric for a fully automated 3-tier Classic LAN deployment with Nexus 9000 and 7000 switches. | Easy_Fabric_Classic | LAN Fabrics, on page 39 |
| Campus VXLAN EVPN | Fabric for a VXLAN EVPN Campus deployment with Catalyst 9000 switches. | Easy_Fabric_IOS_XE | Create Easy_Fabric for Cisco Catalyst 9000 Series Switches, on page 101 |
| BGP Fabric | Fabric for an eBGP based deployment with Nexus 9000 and 3000 switches. Optionally VXLAN EVPN can be enabled on top of the eBGP underlay. | Easy_Fabric_eBGP | Creating VXLAN EVPN Fabric with eBGP-based Underlay, on page 539 |
| Flexible Network | Fabric for flexible deployments with a mix of Nexus and Non-Nexus devices. | External_Fabric | Creating an External Fabric, on page 11 |
| Fabric Group | Domain that can contain Enhanced Classic LAN, Classic LAN, and External Connectivity Network fabrics. | Fabric_Group | LAN Fabrics, on page 39 |
| Classic LAN | Fabric to manage a legacy Classic LAN deployment with Nexus switches. | LAN_Classic | LAN Fabrics, on page 39 |
| LAN Monitor | Fabric for monitoring Nexus switches for basic discovery and inventory management. | LAN_Monitor | LAN Fabrics, on page 39 |
| VXLAN EVPN Multi-Site | Domain that can contain multiple VXLAN EVPN Fabrics (with Layer-2/Layer-3 Overlay Extensions) and other Fabric Types. | MSD_Fabric | Creating the VXLAN EVPN Multi-Sit and Associating Member Fabrics, on page 593 |
| Classic IPFM | Fabric to manage or monitor existing Nexus 9000 switches in an IP Fabric for Media Deployment. | IPFM_Classic | Creating an IPFM Classic Fabric, on page 133 |
| IPFM | Fabric for a fully automated deployment of IP Fabric for Media Network with Nexus 9000 switches. | Easy_Fabric_IPFM | Creating an IPFM Easy Fabric, on pag 136 |
| Multi-Site Interconnect Network | Fabric to interconnect VXLAN EVPN for Multi-Site deployments with a mix of Nexus and Non-Nexus devices | External_Fabric | Creating an External Fabric, on page 11 |

| Type of Fabric | Description | REST API Template Name | Detailed Procedures |
|---|---|---|---|
| External Connectivity Network | Fabric for core and edge router deployments with a mix of Nexus and Non-Nexus devices. | External_Fabric | Creating an External Fabric, on pag |

**Note** Any reference to External_Fabric in this document refers to one of the following 3 fabric templates. Choose the required fabric based on the description in the above table.

- Multi-Site Interconnect Network
- External Connectivity Network
- Flexible Network

The type of fabric will be seen as **External_Fabric** aka the fabric template name, in the following cases:

1. Upgrade and Restore from DCNM 11.5(4) .
2. Upgrade from NDFC 12.0.2f/12.1.1e

All existing functionalists will continue to work similarly to the previous release. You can optionally edit the fabric and choose one of the three options **Flexible Network**, **External Connectivity Network**, **Multi-Site Interconnect Network**. If you edit the fabric settings without choosing one of these options, then the default option **Flexible Network** will be picked. You can toggle between these three options as desired without any loss of functionality. The type of fabric is stored in nvPairs in a variable called **EXT_FABRIC_TYPE**. This can be optionally provided in the payload during fabric creation. If not provided, the default option of **Flexible Network** is picked.

### Prerequisites to Creating a Fabric

- Update the ESXi host settings in the vSphere Client to accept overriding changes in promiscuous mode. For more information, see the *Overriding the Changes in Promiscuous Mode* section.
- Configure the persistent IP addresses in Cisco Nexus Dashboard. For more information, see *Cluster Configuration* section in Cisco Nexus Dashboard User Guide.

## Changing Persistent IP Address

From Cisco NDFC Release 12.1.2e, you can change persistent IP addresses which are assigned for mandatory pods such as POAP-SCP and SNMP trap. To change the persistent IP address, perform the following steps:

### Procedure

**Step 1** On Cisco NDFC Web UI, navigate to **Settings** > **Server Settings** > **Admin** under **LAN Device Management Connectivity** drop-down list change **Management** to **Data** or conversely.

Changing option results in migration of SNMP and POAP-SCP pods to the persistent IP addresses associated with **External Service Pool** on Nexus Dashboard connected with the new **LAN Device Management Connectivity** option. After the completion of this process, the following message is displayed:

Some features have been updated. Reload the page to see latest changes.

Click **Reload the page**.

**Step 2** On Cisco Nexus Dashboard Web UI, navigate to **Infrastructure** > **Cluster Configuration** > **General**, in **External Service Pools** card, change the required IP addresses for **Management Service IP Usage** or **Data Service IP Usage**.

**Step 3** Navigate to NDFC Web UI **Server Settings** page, change the option in **LAN Device Management Connectivity** drop-down list to its initial selection.

Restoring this option to initial settings, results in migration of the SNMP and POAP-SCP pods to use the updated persistent IP address from the appropriate External Service IP pool.

# ESXi Networking for Promiscuous Mode

From Cisco NDFC Release , you can run NDFC on top of virtual Nexus Dashboard (vND) instance with promiscuous mode that is disabled on port groups that are associated with Nexus Dashboard interfaces where External Service IP addresses are specified. vND comprises Nexus Dashboard management interface and data interface. By default, for fabric controller persona, two external service IP addresses are required for the Nexus Dashboard management interface subnet.

Before the NDFC Release , if Inband management or Endpoint Locator or POAP feature was enabled on NDFC, you must also enable promiscuous mode for the Nexus Dashboard data or fabric interface port-group. This setting was mandatory for traffic flow that is associated for these features.

Enabling promiscuous mode raise risk of security issues in NDFC, it is recommended to set default setting for promiscuous mode.

**Note**

- Disabling promiscuous mode is supported from Cisco Nexus Dashboard Release .
- You can disable promiscuous mode when Nexus Dashboard nodes are layer-3 adjacent on the Data network, BGP is configured, and fabric switches are reachable through the data interface.
- You can disable promiscuous mode when Nexus Dashboard interfaces are layer-2 adjacent to switch mgmt0 interface.

If Inband management or EPL is enabled, you must specify External Service IP addresses in the Nexus Dashboard data interface subnet. You can disable promiscuous mode for the Nexus Dashboard data or fabric interface port-group. For more information, refer to Cisco Nexus Dashboard Deployment GuideCisco Nexus Dashboard Deployment Guide

**Note** Default option for promiscuous mode is **Reject**.

**Procedure**

**Step 1** Log into your **vSphere** Client.

**Step 2** Navigate to the ESXi host.

**Step 3** Right-click the host and choose **Settings**.

A sub-menu appears.

**Step 4** Choose **Networking > Virtual Switches**.

All the virtual switches appear as blocks.

**Step 5** Click **Edit Settings** of the VM Network.

**Step 6** Navigate to the **Security** tab.

**Step 7** Update the **Promiscuous mode** settings as follows:

- Check the **Override** check box.
- Choose **Accept** from the drop-down list.

**Step 8** Click **OK**.

# Create a Fabric

To create a Fabric using Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN** > **Fabrics**.

**Step 2** From the **Actions** drop-down list, select **Create Fabric**.

**Step 3** Enter the fabric name and click **Choose Template**.

**Step 4** Specify the values for the fabric settings and click **Save**.

# VXLAN EVPN Fabrics Provisioning

Cisco Nexus Dashboard Fabric Controller provides an enhanced “Easy” fabric workflow for unified underlay and overlay provisioning of the 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 you to tailor the fabric to your 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 Nexus Dashboard Fabric Controller.

These devices are placed in a special fabric called the External Fabric. The same Nexus Dashboard Fabric 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.

The Nexus Dashboard Fabric Controller GUI functions for creating and deploying VXLAN BGP EVPN fabrics are as follows:

**LAN > Fabrics > LAN Fabrics Create Fabric** under **Actions** drop-down list.

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]).

**LAN > Interfaces > LAN Fabrics Create New Interface** under **Actions** drop-down list.

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.

**LAN > Switches > LAN Fabrics Add** under **Actions** drop-down list.

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 Nexus Dashboard Fabric Controller.

**LAN > Services** menu option.

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 Fabric Controller, is covered in the Creating Networks and Creating VRFs in the Networks and VRFs sections.

### Guidelines for VXLAN BGP EVPN Fabrics Provisioning

- For any switch to be successfully imported into Nexus Dashboard Fabric Controller, 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.
  - Ability to execute the **guestshell** commands, which are prefixed by **run guestshell** for the Nexus Dashboard Fabric Controller tracker.
- 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 Nexus Dashboard Fabric Controller 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 clean up 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 Nexus Dashboard Fabric Controller, 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, Nexus Dashboard Fabric Controller moves this device to a migration mode. Once the user sets the appropriate LAN credentials for that device, a subsequent Save & Deploy retriggers 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 **Recalculate Config**, the changes in the fabric are evaluated, and the configuration for the entire fabric is generated. Click **Preview Config** to preview the generated configuration, and then deploy it at a fabric level. Therefore, **Deploy Config** 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 to switches** option to deploy per switch configurations. This option is a local operation for a switch, that is, the expected configuration or intent for a switch is evaluated against it's current running configuration, and a config compliance check is performed for the switch to get the **In-Sync** or **Out-of-Sync** status. If the switch is out of sync, the user is provided with a preview of all the configurations running in that particular switch that vary from the intent defined by the user for that respective switch.

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

  The intent in Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller 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 doesn't match the policies in Nexus Dashboard Fabric Controller, 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 VXLAN EVPN Fabric Using the Easy_Fabric Template

This topic describes how to create a new VXLAN EVPN fabric using the **Easy_Fabric** template and contains descriptions for the IPv4 underlay. For information about the IPv6 underlay, see IPv6 Underlay Support for Easy Fabric, on page 69.

1. Navigate to the **LAN Fabrics** page:

   **LAN** > **Fabrics**

2. Click **Actions** > **Create Fabric**.

   The **Create Fabric** window appears.

3. Enter a unique name for the fabric in the **Fabric Name** field, then click **Choose Fabric**.

   A list of all available fabric templates are listed.

4. From the available list of fabric templates, choose the **Easy_Fabric** template, then click **Select**.

5. Enter the necessary field values to create a fabric.

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



**Note** If you're 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), see Multi-Site Domain for VXLAN BGP EVPN Fabrics , on page 589 before creating the member fabric.

- General Parameters, on page 50
- Replication, on page 51
- VPC, on page 52
- Protocols, on page 53
- Advanced, on page 57
- Resources, on page 61
- Manageability, on page 64
- Bootstrap, on page 65
- Configuration Backup, on page 66
- Flow Monitor, on page 67

6. When you have completed the necessary configurations, click **Save**.
   - Click on the fabric to display a summary in the slide-in pane.
   - Click on the Launch icon to display the Fabric Overview.

### General Parameters

The **General Parameters** tab is displayed by default. The fields in this tab are described in the following table.

| Field | Description |
|---|---|
| **BGP ASN** | Enter the BGP AS number the fabric is associated with. This must be same as existing fabric. |
| **Enable IPv6 Underlay** | Enable the IPv6 underlay feature. For information, see IPv6 Underlay Support for Easy Fabric, on page 69. |
| **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 Subnet IPv6 Mask** | Specifies the subnet mask for the fabric interface IPv6 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.<br><br>To deploy spine devices as RRs, Nexus Dashboard Fabric Controller 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 won't change.<br><br>*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.<br><br>*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.<br><br>**1.** Change the value in the drop-down box to 2.<br><br>**2.** Identify the spine switches designated as route reflectors.<br><br>An instance of the **rr_state** policy is applied on the spine switch if it's 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.<br><br>**3.** Delete the unneeded spine devices from the fabric (right-click the spine switch icon and choose **Discovery > Remove from fabric**).<br><br>If you delete existing RR devices, the next available spine switch is selected as the replacement RR.<br><br>**4.** Click **Deploy Config** in the fabric topology window.<br><br>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. |

| Field | Description |
|---|---|
| **Enable Performance Monitoring** | Check the check box to enable performance monitoring.<br>Ensure that you do not clear interface counters from the Command Line Interface of the switches. Clearing interface counters can cause the Performance Monitor to display incorrect data for traffic utilization. If you must clear the counters and the switch has both `clear counters` and `clear counters snmp` commands (not all switches have the `clear counters snmp` command), ensure that you run both the main and the SNMP commands simultaneously. For example, you must run the `clear counters interface ethernet` *`slot/port`* command followed by the `clear counters interface ethernet` *`slot/port`* `snmp` command. This can lead to a one time spike. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Replication

The fields in the **Replication** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **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.<br>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.<br>The replication mode change isn't allowed if a policy template instance is created for the current mode. For example, if a multicast related policy is created and deployed, you can't change the mode to Ingress. |
| **Enable Tenant Routed Multicast (TRM)** | Check 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**.<br>For more information, see Overview of Tenant Routed Multicast, on page 69. |
| **Rendezvous-Points** | Enter the number of spine switches acting as rendezvous points. |

| Field | Description |
|---|---|
| **RP mode** | Choose from the two supported multicast modes of replication, ASM (for Any-Source Multicast [ASM]) or BiDir (for Bidirectional PIM [BIDIR-PIM]).<br>When you choose ASM, the BiDir related fields aren't enabled. When you choose BiDir, the BiDir related fields are enabled.<br>**Note**<br>BIDIR-PIM is supported on Cisco's Cloud Scale Family platforms 9300-EX and 9300-FX/FX2, and software release 9.2(1) onwards.<br>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. |
| **Underlay Primary RP Loopback ID** | Enabled if you choose BIDIR-PIM as the multicast mode of replication.<br>The primary loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay. |
| **Underlay Backup RP Loopback ID** | Enabled if you choose BIDIR-PIM as the multicast mode of replication.<br>The secondary loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay. |
| **Underlay Second Backup RP Loopback Id** | Used for the second fallback Bidir-PIM Phantom RP. |
| **Underlay Third Backup RP Loopback Id** | Used for the third fallback Bidir-PIM Phantom RP. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### VPC

The fields in the **VPC** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **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.<br>If you use IPv6 addresses, you must use loopback IDs. |
| **vPC Auto Recovery Time** | Specifies the vPC auto recovery time-out period in seconds. |

| Field | Description |
|---|---|
| **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. Uncheck the check box to disable the function. |
| **vPC advertise-pip** | Select the check box to enable the Advertise PIP feature.<br>You can enable the advertise PIP feature on a specific vPC as well. . |
| **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. .<br>**Note**<br>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**. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Protocols

The fields in the **Protocols** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **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 Anycast Loopback Id** | The loopback interface ID is greyed out and used for vPC Peering in VXLANv6 Fabrics only. |
| **Underlay Routing Protocol Tag** | The tag defining the type of network. |

| Field | Description |
|---|---|
| **OSPF Area ID** | The OSPF area ID, if OSPF is used as the IGP within the fabric.<br>**Note**<br>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.<br>**Note**<br>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 Network Point-to-Point** | Enables network point-to-point on fabric interfaces which are numbered. |
| **Enable IS-IS Authentication** | Select the check box to enable IS-IS authentication. Deselect the check box to disable it. If you enable this field, the IS-IS authentication fields are enabled. |
| **IS-IS Authentication Keychain Name** | Enter the Keychain name, such as CiscoisisAuth. |
| **IS-IS Authentication Key ID** | The Key ID is populated. |
| **IS-IS Authentication Key** | Enter the Cisco Type 7 encrypted key.<br>**Note**<br>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. |
| **Set IS-IS Overload Bit** | When enabled, set the overload bit for an elapsed time after a reload. |
| **IS-IS Overload Bit Elapsed Time** | Allows you to clear the overload bit after an elapsed time in seconds. |
| **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.<br>**Note**<br>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. |

| Field | Description |
|---|---|
| **BGP Authentication Key** | Enter the encrypted key based on the encryption type.<br>**Note**<br>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** | Select this check box to enable PIM hello authentication on all the intra-fabric interfaces of the switches in a fabric. This check box is editable only for the Multicast replication mode. Note this check box is valid only for the IPv4 underlay. |
| **PIM Hello Authentication Key** | Specifies the PIM hello authentication key. For more information, see Retrieving PIM Hello Authentication Key.<br>To retrieve the PIM Hello Authentication Key, perform the following steps:<br>**1.** SSH into the switch.<br>**2.** On an unused switch interface, enable the following:<br>`switch(config)# interface e1/32`<br>`switch(config-if)# ip pim hello-authentication ah-md5 pimHelloPassword`<br>In this example, **pimHelloPassword** is the cleartext password that has been used.<br>**3.** Enter the **show run interface** command to retrieve the PIM hello authentication key.<br>`switch(config-if)# show run interface e1/32 \| grep pim`<br>`ip pim sparse-mode`<br>`ip pim hello-authentication ah-md5 3 d34e6c5abc7fecf1caa3b588b09078e0`<br>In this example, **d34e6c5abc7fecf1caa3b588b09078e0** is the PIM hello authentication key that should be specified in the fabric settings. |
| **Enable BFD** | Check 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.<br>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.<br>The following config is pushed after you select the **Enable BFD** check box:<br>`feature bfd`<br>For information about BFD feature compatibility, refer your respective platform documentation and for information about the supported software images, see *Compatibility Matrix for Cisco Nexus Dashboard Fabric Controller*. |
| **Enable BFD for iBGP** | Check the check box to enable BFD for the iBGP neighbor. This option is disabled by default. |
| **Enable BFD for OSPF** | Check 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** | Check 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. |

| Field | Description |
|---|---|
| **Enable BFD for PIM** | Check 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.<br>Following are examples of the BFD global policies:<br>`router ospf <ospf tag>`<br>`  bfd`<br>`router isis <isis tag>`<br>`  address-family ipv4 unicast`<br>`  bfd`<br>`ip pim bfd`<br>`router bgp <bgp asn>`<br>`  neighbor <neighbor ip>`<br>`  bfd` |
| **Enable BFD Authentication** | Check the check box to enable BFD authentication. If you enable this field, the **BFD Authentication Key ID** and **BFD Authentication Key** fields are editable.<br>**Note**<br>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.<br>For information about how to retrieve the BFD authentication parameters. . |

| Field | Description |
|---|---|
| **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.<br>If you use BGP templates, add the authentication configuration within the template and uncheck the Enable BGP Authentication check box to avoid duplicate configuration.<br>In the sample configuration, the 3DES password is displayed after password 3.<br>`router bgp 65000`<br>`  password 3 sd8478fswerdfw3434fsw4f4w34sdsd8478fswerdfw3434fsw4f4w`<br>The following fields can be used to specify different configurations:<br>• **iBGP Peer-Template Config** – Specifies the config used for RR and spines with border role.<br>• **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).<br>In a 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. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Advanced

The fields in the **Advanced** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **VRF Template** | Specifies the VRF template for creating VRFs. |
| **Network Template** | Specifies the network template for creating networks. |
| **VRF Extension Template** | Specifies the VRF extension template for enabling VRF extension to other fabrics. |
| **Network Extension Template** | Specifies the network extension template for extending networks to other fabrics. |
| **Overlay Mode** | VRF/Network configuration using config-profile or CLI, default is config-profile. For more information, see Overlay Mode, on page 83. |
| **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. |

| Field | Description |
|---|---|
| **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. |
| **Unshut Host Interfaces by Default** | Check this check box to unshut the host interfaces by default. |
| **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.<br>The following list describes the variables in the syntax:<br>• **$$VNI$$**: Specifies the network VNI ID found in the switch configuration. This is a mandatory keyword required to create unique network names.<br>• **$$VLAN_ID$$**: Specifies the VLAN ID associated with the network.<br>VLAN ID is specific to switches, hence Nexus Dashboard Fabric Controller picks the VLAN ID from one of the switches, where the network is found, randomly and use it in the name.<br>We recommend not to use this unless the VLAN ID is consistent across the fabric for the VNI.<br>• **<string>**: This variable is optional and you can enter any number of alphanumeric characters that meet the network name guidelines.<br>An example overlay network name: Site_VNI12345_VLAN1234<br>**Note**<br>Ignore this field for greenfield deployments. The **Brownfield Overlay Network Name Format** applies for the following brownfield imports:<br>• CLI-based overlays<br>• Configuration profile-based overlay |
| **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. |

| Field | Description |
|---|---|
| **Enable VXLAN OAM** | Enables the VXLAM OAM functionality for devices in the fabric. This is enabled by default. Uncheck the check box to disable VXLAN OAM function.<br>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.<br>**Note**<br>The VXLAN OAM feature in Cisco Nexus Dashboard Fabric Controller is only supported on a single fabric or site. |
| **Enable Tenant DHCP** | Check 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.<br>**Note**<br>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 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 Nexus Dashboard Fabric Controller, 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.<br>**Note**<br>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)** | Check 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. It enables bi-directional compliance checks to flag additional configs in the running config that are not in the intent/expected config. By default, this feature is disabled. |
| **Enable AAA IP Authorization** | Enables AAA IP authorization, when IP Authorization is enabled in the remote authentication server. This is required to support Nexus Dashboard Fabric Controller in scenarios where customers have strict control of which IP addresses can have access to the switches. |
| **Enable NDFC as Trap Host** | Select this check box to enable Nexus Dashboard Fabric Controller as an SNMP trap destination. Typically, for a native HA Nexus Dashboard Fabric Controller deployment, the eth1 VIP IP address will be configured as SNMP trap destination on the switches. By default, this check box is enabled. |
| **Anycast Border Gateway advertise-pip** | Enables to advertise Anycast Border Gateway PIP as VTEP. Effective on MSD fabric 'Recalculate Config'. |

| Field | Description |
|---|---|
| **Greenfield Cleanup Option** | Enable the switch cleanup option for switches imported into Nexus Dashboard Fabric Controller 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 check 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 79. |
| **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 Nexus Dashboard Fabric Controller.<br>If the PTP loopback ID is not found during **Deploy Config**, the following error is generated:<br>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** | Check the check box to enable the MPLS Handoff feature. For more information, see the MPLS SR and LDP Handoff, on page 619 chapter in External/WAN Layer 3 Connectivity for VXLAN BGP EVPN Fabrics. |
| **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. 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.<br>Review the actual queuing policies by opening the policy file in the template editor. From Cisco Nexus Dashboard Fabric Controller Web UI, choose **Operations > Templates**. Search for the queuing policies by the policy file name, for example, **queuing_policy_default_8q_cloudscale**. Choose the file. From the **Actions** drop-down list, select **Edit template content** to edit the policy.<br>See the *Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide* for platform specific details. |

| Field | Description |
|---|---|
| **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 Enabling MACsec.<br>*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 should be added via the switch freeform on NDFC. For more information, see Enabling Freeform Configurations on Fabric Switches , on page 95. |
| **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. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Resources

The fields in the **Resources** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **Manual Underlay IP Address Allocation** | *Do not* check this check box if you are transitioning your VXLAN fabric management to Nexus Dashboard Fabric Controller.<br>• By default, Nexus Dashboard Fabric Controller allocates the underlay IP address resources (for loopbacks, fabric interfaces, etc) dynamically from the defined pools. If you check the check box, the allocation scheme switches to static, and some of the dynamic IP address range fields are disabled.<br>• For static allocation, the underlay IP address resources must be populated into the Resource Manager (RM) using REST APIs.<br>• The Underlay RP Loopback IP Range field stays enabled if BIDIR-PIM function is chosen for multicast replication.<br>• 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.<br>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. |
| **Underlay Routing Loopback IPv6 Range** | Specifies Loopback0 IPv6 Address Range |
| **Underlay VTEP Loopback IPv6 Range** | Specifies Loopback1 and Anycast Loopback IPv6 Address Range. |
| **Underlay Subnet IPv6 Range** | Specifies IPv6 Address range to assign Numbered and Peer Link SVI IPs. |
| **BGP Router ID Range for IPv6 Underlay** | Specifies BGP router ID range for IPv6 underlay. |
| **Layer 2 VXLAN VNI Range** | Specifies the overlay VXLAN VNI range for the fabric (min:1, max:16777214). |
| **Layer 3 VXLAN VNI Range** | Specifies the overlay VRF VNI range for the fabric (min:1, max:16777214). |
| **Network VLAN Range** | VLAN range for the per switch overlay network (min:2, max:4094). |

| Field | Description |
|---|---|
| **VRF VLAN Range** | VLAN range for the per switch overlay Layer 3 VRF (min:2, max:4094). |
| **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.<br>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 Back2Back&ToExternal, then VRF Lite IFCs are auto-created. |
| **Auto Deploy for Peer** | This check box is applicable for VRF Lite deployment. When you select this checkbox, auto-created VRF Lite IFCs will have the **Auto Generate Configuration for Peer** field in the **VRF Lite** tab set.<br>To access VRF Lite IFC configuration, navigate to the **Links** tab, select the particular link, and then choose **Actions** > **Edit**.<br>You can check or uncheck the check box when the **VRF Lite Deployment** field is not set to **Manual**. This configuration only affects the new auto-created IFCs and does not affect the existing IFCs. You can edit an auto-created IFC and check or uncheck the **Auto Generate Configuration for Peer** field. This setting takes priority always. |
| **Auto Deploy Default VRF** | When you select this check box, the **Auto Generate Configuration on default VRF** field is automatically enabled for auto-created VRF Lite IFCs. You can check or uncheck this check box when the **VRF Lite Deployment** field is not set to **Manual**. The **Auto Generate Configuration on default VRF** field when set, automatically configures the physical interface for the border device, and establishes an EBGP connection between the border device and the edge device or another border device in a different VXLAN EVPN fabric. |
| **Auto Deploy Default VRF for Peer** | When you select this check box, the **Auto Generate Configuration for NX-OS Peer on default VRF** field is automatically enabled for auto-created VRF Lite IFCs. You can check or uncheck this check box when the **VRF Lite Deployment** field is not set to **Manual**. The **Auto Generate Configuration for NX-OS Peer on default VRF** field when set, automatically configures the physical interface and the EBGP commands for the peer NX-OS switch.<br>**Note**<br>To access the **Auto Generate Configuration on default VRF** and **Auto Generate Configuration for NX-OS Peer on default VRF** fields for an IFC link, navigate to the **Links** tab, select the particular link and choose **Actions** > **Edit**. |
| **Redistribute BGP Route-map Name** | Defines the route map for redistributing the BGP routes in default VRF. |

| Field | Description |
|---|---|
| **VRF Lite Subnet IP Range** and **VRF Lite Subnet Mask** | These fields are populated with the DCI subnet details. Update the fields as needed.<br>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:<br>**Note**<br>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.<br>**1.** Update the L2 range and click **Save**.<br>**2.** 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. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Manageability

The fields in the **Manageability** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **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.<br>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. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Bootstrap

The fields in the **Bootstrap** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **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.<br>After you enable bootstrap, you can enable the DHCP server for automatic IP address assignment using one of the following methods:<br>• External DHCP Server: Enter information about the external DHCP server in the **Switch Mgmt Default Gateway** and **Switch Mgmt IP Subnet Prefix** fields.<br>• 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.<br>If you do not select this check box, Nexus Dashboard Fabric Controller 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.<br>**Note**<br>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.<br>*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. |

| Field | Description |
|---|---|
| **Enable AAA Config** | Select this check box to include AAA configurations from the Manageability tab as part of the device start-up config post bootstrap. |
| **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.<br>The format of the scope should be defined as:<br>**DHCP Scope Start Address, DHCP Scope End Address, Switch Management Default Gateway, Switch Management Subnet Prefix**<br>For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24 |
| **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.<br>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 Enabling Freeform Configurations on Fabric Switches , on page 95. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Configuration Backup

The fields in the **Configuration Backup** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **Hourly Fabric Backup** | Select the check box to enable an hourly backup of fabric configurations and the intent.<br>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. |

| Field | Description |
|---|---|
| **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.<br><br>Select both the check boxes to enable both back up processes.<br><br>The backup process is initiated after you click **Save**.<br><br>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.<br><br>The number of fabric backups that will be retained on NDFC is decided by the **Settings** > **Server Settings** > **LAN Fabric** > **Maximum Backups per Fabric**.<br><br>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.<br><br>**Note**<br>To trigger an immediate backup, do the following:<br><br>**1.** Choose **LAN > Topology**.<br><br>**2.** Click within the specific fabric box. The fabric topology screen comes up.<br><br>**3.** From the **Actions** pane at the left part of the screen, click **Re-Sync Fabric**.<br><br>You can also initiate the fabric backup in the fabric topology window. Click **Backup Now** in the **Actions** pane. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Flow Monitor

The fields in the **Flow Monitor** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **Enable Netflow** | Check this check box to enable Netflow on VTEPs for this Fabric. By default, Netflow is disabled. On Enable, NetFlow configuration will be applied to all VTEPS that support netflow.<br><br>**Note**<br>When Netflow is enabled on the fabric, you can choose not to have netflow on a particular switch by having a dummy no_netflow PTI.<br><br>If netflow is not enabled at the fabric level, an error message is generated when you enable netflow at the interface, network, or vrf level. For information about Netflow support for Cisco NDFC, refer to Netflow Support, on page 149. |

In the **Netflow Exporter** area, click **Actions > Add** to add one or more Netflow exporters. This exporter is the receiver of the netflow data. The fields on this screen are:

- **Exporter Name** – Specifies the name of the exporter.

- **IP** – Specifies the IP address of the exporter.
- **VRF** – Specifies the VRF over which the exporter is routed.
- **Source Interface** – Enter the source interface name.
- **UDP Port** – Specifies the UDP port over which the netflow data is exported.

Click **Save** to configure the exporter. Click **Cancel** to discard. You can also choose an existing exporter and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

In the **Netflow Record** area, click **Actions > Add** to add one or more Netflow records. The fields on this screen are:

- **Record Name** – Specifies the name of the record.
- **Record Template** – Specifies the template for the record. Enter one of the record templates names. In Release 12.0.2, the following two record templates are available for use. You can create custom netflow record templates. Custom record templates saved in the template library are available for use here.
  - **netflow_ipv4_record** – to use the IPv4 record template.
  - **netflow_l2_record** – to use the Layer 2 record template.
- **Is Layer2 Record** – Check this check box if the record is for Layer2 netflow.

Click **Save** to configure the report. Click **Cancel** to discard. You can also choose an existing record and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

In the **Netflow Monitor** area, click **Actions > Add** to add one or more Netflow monitors. The fields on this screen are:

- **Monitor Name** – Specifies the name of the monitor.
- **Record Name** – Specifies the name of the record for the monitor.
- **Exporter1 Name** – Specifies the name of the exporter for the netflow monitor.
- **Exporter2 Name** – (optional) Specifies the name of the secondary exporter for the netflow monitor.

The record name and exporters referred to in each netflow monitor must be defined in "**Netflow Record**" and "**Netflow Exporter**".

Click **Save** to configure the monitor. Click **Cancel** to discard. You can also choose an existing monitor and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

# Configuring Fabrics with eBGP Underlay

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

# IPv6 Underlay Support for Easy Fabric

You can create an Easy fabric with IPv6 only underlay. The IPv6 underlay is supported only for the **Easy_Fabric** template. For more information, see *Configuring a VXLANv6 Fabric*.

# Overview of Tenant Routed Multicast

Tenant Routed Multicast (TRM) enables multicast forwarding on the VXLAN fabric that uses a BGP-based EVPN control plane. TRM provides multi-tenancy aware multicast forwarding between senders and receivers within the same or different subnet local or across VTEPs.

With TRM enabled, multicast forwarding in the underlay is leveraged to replicate VXLAN encapsulated routed multicast traffic. A Default Multicast Distribution Tree (Default-MDT) is built per-VRF. This is an addition to the existing multicast groups for Layer-2 VNI Broadcast, Unknown Unicast, and Layer-2 multicast replication group. The individual multicast group addresses in the overlay are mapped to the respective underlay multicast address for replication and transport. The advantage of using a BGP-based approach allows the VXLAN BGP EVPN fabric with TRM to operate as fully distributed Overlay Rendezvous-Point (RP), with the RP presence on every edge-device (VTEP).

A multicast-enabled data center fabric is typically part of an overall multicast network. Multicast sources, receivers, and multicast rendezvous points might reside inside the data center but also might be inside the campus or externally reachable via the WAN. TRM allows a seamless integration with existing multicast networks. It can leverage multicast rendezvous points external to the fabric. Furthermore, TRM allows for tenant-aware external connectivity using Layer-3 physical interfaces or subinterfaces.

For more information, see the following:

- Guidelines and Limitations for Tenant Routed Multicast
- Guidelines and Limitations for Layer 3 Tenant Routed Multicast
- Guidelines and Limitations for Layer 2/Layer 3 Tenant Routed Multicast (Mixed Mode)

## Overview of Tenant Routed Multicast with VXLAN EVPN Multi-Site

Tenant Routed Multicast with Multi-Site enables multicast forwarding across multiple VXLAN EVPN fabrics connected via Multi-Site.

The following two use cases are supported:

- Use Case 1: TRM provides Layer 2 and Layer 3 multicast services across sites for sources and receivers across different sites.
- Use Case 2: Extending TRM functionality from VXLAN fabric to sources receivers external to the fabric.

TRM Multi-Site is an extension of BGP-based TRM solution that enables multiple TRM sites with multiple VTEPs to connect to each other to provide multicast services across sites in most efficient possible way. Each TRM site is operating independently and border gateway on each site allows stitching across each site. There can be multiple Border Gateways for each site. In a given site, the BGW peers with Route Sever or BGWs of other sites to exchange EVPN and MVPN routes. On the BGW, BGP will import routes into the local VRF/L3VNI/L2VNI and then advertise those imported routes into the Fabric or WAN depending on where the routes were learnt from.

# Tenant Routed Multicast with VXLAN EVPN Multi-Site Operations

The operations for TRM with VXLAN EVPN Multi-Site are as follows:

- Each Site is represented by Anycast VTEP BGWs. DF election across BGWs ensures no packet duplication.
- Traffic between Border Gateways uses ingress replication mechanism. Traffic is encapsulated with VXLAN header followed by IP header.
- Each Site will only receive one copy of the packet.
- Multicast source and receiver information across sites is propagated by BGP protocol on the Border Gateways configured with TRM.
- BGW on each site receives the multicast packet and re-encapsulate the packet before sending it to the local site.

For information about guidelines and limitations for TRM with VXLAN EVPN Multi-Site, see Configuring Tenant Routed Multicast.

# Configuring TRM for Single Site Using Cisco Nexus Dashboard Fabric Controller

This section is assumes that a VXLAN EVPN fabric has already been provisioned using Cisco Nexus Dashboard Fabric Controller.

### Procedure

**Step 1** Enable TRM for the selected Easy Fabric. If the fabric template is **Easy_Fabric**, from the Fabric Overview **Actions** drop-down, choose the **Edit Fabric** option. Click the **Replication** tab. The fields on this tab are:

**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**: When you select the **Enable Tenant Routed Multicast (TRM)** check box, the multicast address for Tenant Routed Multicast traffic is auto 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. From the Fabric Overview **Actions** drop-down list, choose **Recalculate Config** and then choose **Deploy Config** 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 **EVPN** tab.

**Step 2** Enable TRM for the VRF.

Navigate to **Fabric Overview > VRFs > 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 Mcast Address** – Specifies the multicast address associated with the VRF. The multicast address is used for transporting multicast traffic in the fabric underlay.

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

Click **Save** to save the settings. The switches go into the pending state, that is, blue color. These settings enable the following:

- Enable PIM on L3VNI SVI.
- Route-Target Import and Export for MVPN AFI.
- RP and other multicast configuration for the VRF.
- Loopback interface using the above RP address and RP loopback id for the distributed RP.

**Step 3** Enable TRM for the network.

Navigate to **Fabric Overview > Networks > Networks**. Edit the selected network and navigate to the **Advanced** tab. Edit the following TRM setting:

**TRM Enable** – Select the check box to enable TRM.

Click **Save** to save the settings. The switches go into the pending state, that is, the blue color. The TRM settings enable the following:

- Enable PIM on the L2VNI SVI.
- Create a PIM policy **none** to avoid PIM neighborship with PIM Routers within a VLAN. The **none** keyword is a configured route map to deny any ipv4 addresses to avoid establishing PIM neighborship policy using anycast IP.

# Configuring TRM for Multi-Site Using Cisco Nexus Dashboard Fabric Controller

This section assumes that a Multi-Site Domain (MSD) has already been deployed by Cisco Nexus Dashboard Fabric Controller and TRM needs to be enabled.

### Procedure

**Step 1** Enable TRM on the BGWs.

Navigate to **Fabric Overview > VRFs > 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 Mcast Address** – Specifies the multicast address associated with the VRF. The multicast address is used for transporting multicast traffic in the fabric underlay.

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

**Enable TRM BGW MSite** - Select the check box to enable TRM on Border Gateway Multi-Site.

Click on **Save** to save the settings. The switches go into the pending state, that is, blue color. These settings enable the following:

- Enable feature ngmvpn: Enables the Next-Generation Multicast VPN (ngMVPN) control plane for BGP peering.
- Enables PIM on L3VNI SVI.
- Configures L3VNI Multicast Address.
- Route-Target Import and Export for MVPN AFI.
- RP and other multicast configuration for the VRF.
- Loopback interface for the distributed RP.
- Enable Multi-Site BUM ingress replication method for extending the Layer 2 VNI

**Step 2** Establish MVPN AFI between the BGWs.

Double-click the MSD fabric to open the **Fabric Overview** window. Choose **Links**. Filter it by the policy - **Overlays**.

Select and edit each overlay peering to enable TRM by checking the **Enable TRM** check box.

Click **Save** to save the settings. The switches go into the pending state, that is, the blue color. The TRM settings enable the MVPN peering's between the BGWs, or BGWs and Route Server.

# vPC Fabric Peering

vPC Fabric Peering provides an enhanced dual-homing access solution without the overhead of wasting physical ports for vPC Peer Link. This feature preserves all the characteristics of a traditional vPC. For more information, see *Information about vPC Fabric Peering* section in *Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide*.

You can create a virtual peer link for two switches or change the existing physical peer link to a virtual peer link. Cisco NDFC support vPC fabric peering in both greenfield as well as brownfield deployments. This feature is applicable for **Easy_Fabric** 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 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.
- 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/GX/GX2 platform switches support vPC Fabric Peering. Cisco Nexus 9200 and 9500 platform switches do not support vPC Fabric Peering. For more information, see *Guidelines and Limitations for vPC Fabric Peering* section in *Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide*.
- If you use other Cisco Nexus 9000 Series Switches, a warning will appear during **Recalculate & 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 **Recalculate & Deploy** if you try to pair any of these.
- Brownfield deployments and greenfield deployments support vPC fabric peering in Cisco NDFC.

- However, you can import switches that are connected using physical peer links and convert the physical peer links to virtual peer links after **Recalculate & Deploy**. To update a TCAM region during the feature configuration, use the **hardware access-list tcam ingress-flow redirect** *512* command in the configuration terminal.

### QoS for Fabric vPC-Peering

In the **Easy_Fabric** 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 Creating a VXLAN EVPN Fabric Using the Easy_Fabric Template, on page 49.

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 configuration 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 configuration should make sure to honor the brownfield configuration.
- 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 configuration.
- 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 **Recalculate & 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 Creating a VXLAN EVPN Fabric Using the Easy_Fabric Template, on page 49.

### 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.<br>**Note**<br>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 NDFC Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN** > **Fabrics**.

The **LAN Fabrics** window appears.

**Step 2** Choose a fabric with the **Easy_Fabric** or **Easy_Fabric_eBGP**fabric templates.

**Step 3** On the **Topology** window, 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 **Fabric Overview** window. Choose a switch in the **Switches** tab and click on **Actions** > **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 **Recalculate & Deploy**.

**Step 6** Click **Save**.

**Step 7** In the **Topology** window, choose **Recalculate & Deploy**.

The **Deploy Configuration** 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 side-by-side configuration.

**Step 10** Close the window.

**Step 11** Click the pending errors icon next to **Recalculate & 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 the topology window. For more information, see *Guidelines and Limitations for vPC Fabric Peering* and *Migrating from vPC to vPC Fabric Peering* sections in *Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide*.

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 NDFC Web UI, perform the following steps:

### Before you begin

- Perform 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.
  - Cisco Nexus 9300-EX, and 9300-FX/FXP/FX2/FX3/GX/GX2 platform switches. For more information, see *Guidelines and Limitations for vPC Fabric Peering* section in *Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide*.

### Procedure

**Step 1** Choose **LAN** > **Fabrics**.

The **LAN Fabrics** window appears.

**Step 2** Choose a fabric with the **Easy_Fabric** or **Easy_Fabric_eBGP** fabric templates.

**Step 3** On the **Topology** window, 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 **Fabric Overview** window. Choose a switch in the **Switches** tab and click on **Actions** > **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 **Recalculate & 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 **Topology** window, choose **Recalculate & Deploy**.

The **Deploy Configuration** 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 **Recalculate & 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 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 NDFC 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 **LAN** > **Fabrics**.

The **LAN Fabrics** window appears.

**Step 2** Choose a fabric with the **Easy_Fabric** or **Easy_Fabric_eBGP**fabric templates.

**Step 3** On the **Topology** window, 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 **Fabric Overview** window. Choose a switch in the **Switches** tab and click on **Actions** > **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 **Topology** window, choose **Recalculate & Deploy**.

The **Deploy Configuration** 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 **Recalculate & 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 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.

# Precision Time Protocol for Easy Fabric

In the fabric settings for the **Easy_Fabric** 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 Nexus Dashboard Insights User Guide*.

For Nexus Dashboard Fabric Controller 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 **Deploy Config**. 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 Switch Role

Super Spine is a device that is used for interconnecting multiple spine-leaf PODs. 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 switch roles are supported in NDFC:

- 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 is 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 NDFC:

- Supported with Easy Fabrics only.
- 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.

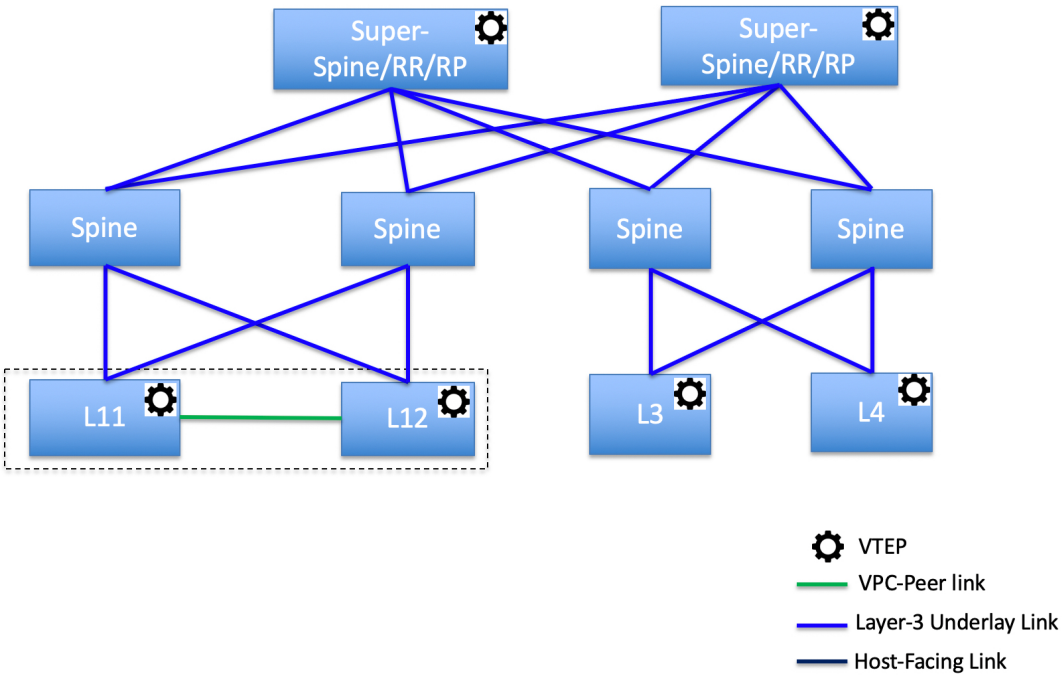
- RR or RP (if applicable) functionality is always be configured on super spines if they are present in a fabric. The maximum number of 4 RRs and RPs are supported even with Super Spines.
- Border Super Spine and Border Gateway Super Spine roles are supported for inter-fabric connections.
- vPC configurations 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

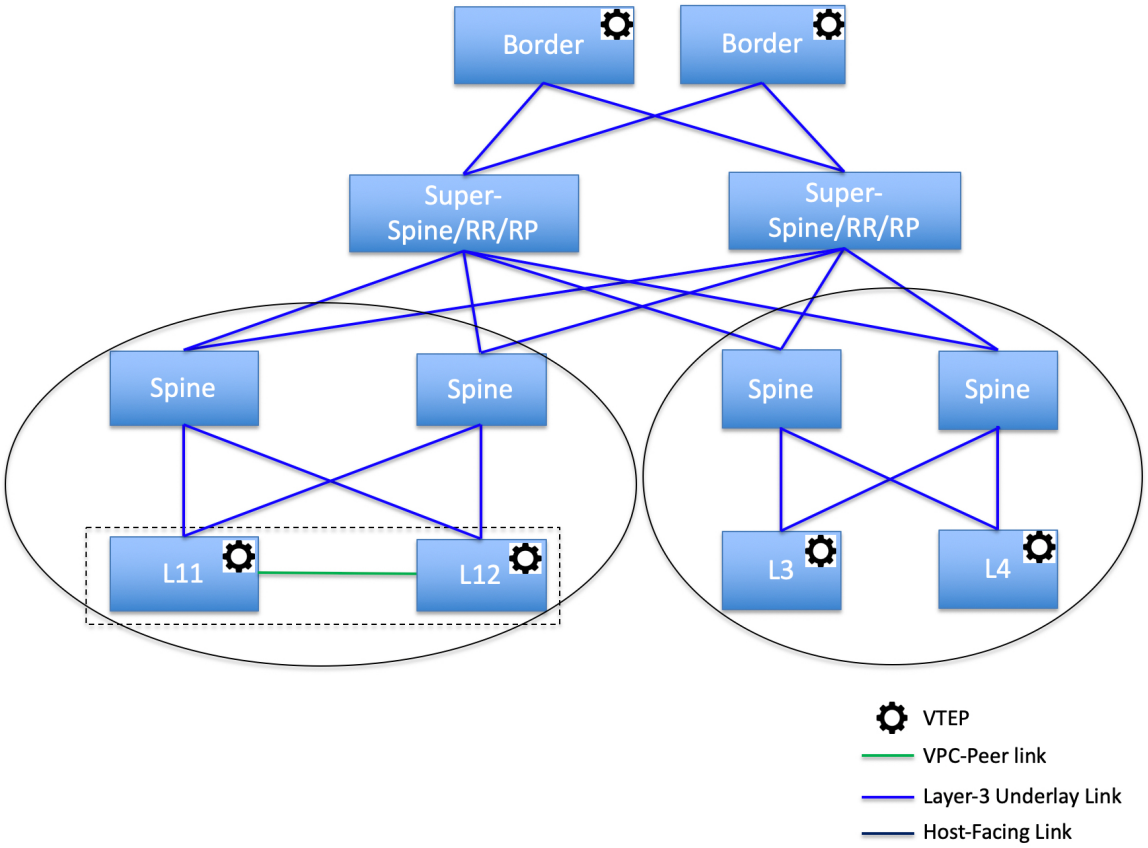
NDFC supports the following topologies with super spine switches.

**Topology 1: Super Spine Switches in a Spine Leaf Topology**




In this topology, leaf switches are connected to spines, and spines are connected to super spine switches which can be super spines, border super spines, and 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 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

To add a super spine switch to an existing VXLAN BGP EVPN fabric, perform the following steps:

### Procedure

**Step 1** Choose **LAN** > **Fabrics**. Double-click on the required fabric.

The **Fabric Overview** window appears.

**Step 2** On the **Switches** tab, click **Actions** > **Add Switches**.

For more information, see Adding Switches to a Fabric, on page 279.

**Step 3** Right-click on 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 exists in the fabric, you can assign border super spine and border gateway super spine roles for any new device.

- If super spine or any of its variation role is not assigned, you may assign the role to any new device if it is connected to a non-border spine. After a **Recalculate & Deploy**, you will receive an error that can be resolved by clicking on the **Resolve** button as shown in the below steps.

**Step 4** On the **Fabric Overview** window, click **Actions** > **Recalculate & Deploy**.

The following error message is displayed:

Super Spine role cannot be allowed in the existing fabric as it is disruptive. Please go to 'Event Analytics' and click on the resolve button to proceed.

**Step 5** Choose **Event Analytics** > **Alarms**, click on the **ID**.

The **Alarm ID** slide-in pane appears.

**Step 6** Click **Resolve**.

The **Confirm action** dialog box appears.

**Step 7** Click **Confirm**.

**Step 8** On the **Fabric Overview** window, click **Actions** > **Recalculate & Deploy**.

Do not add a devices with super spine, border super spine, or border gateway super spine role if the device is connected to a border spine or border gateway spine. This action results in an error after you recalculate and deploy the configuration. To use existing devices with border spine roles, remove the device and add the device with appropriate roles.

# Overlay Mode

You can create a VRF or network in CLI or config-profile mode at the fabric level. The overlay mode of member fabrics of an MSD fabric is set individually at the member-fabric level. Overlay mode can only be changed before deploying overlay configurations to the switches. After the overlay configuration is deployed, you cannot change the mode unless all the VRF and network attachments are removed.



**Note** If you upgrade from Cisco DCNM Release 11.5(x), the existing config-profile mode functions the same.

If the switch has config-profile based overlays, you can import it in the **config-profile** overlay mode only. If you import it in the **cli** overlay mode, an error appears during brownfield import.

For brownfield import, if overlay is deployed as **config-profile** mode, it can be imported in **config-profile** mode only. However, if overlay is deployed as **cli**, it can be imported in either **config-profile** or **cli** modes.

To choose the overlay mode of VRFs or networks in a fabric, perform the following steps:

1. Navigate to the **Edit Fabric** window.
2. Go to the **Advanced** tab.
3. From the **Overlay Mode** drop-down list, choose **config-profile** or **cli**.

   The default mode is **config-profile**.

# Sync up Out-of-Band Switch Interface Configurations

Any interface level configuration made outside of Nexus Dashboard Fabric Controller (via CLI) can be synced to Nexus Dashboard Fabric Controller and then managed from Nexus Dashboard Fabric Controller. Also, the vPC pair configurations are automatically detected and paired. This applies to the External_Fabric and LAN_Classic fabrics only. The vPC pairing is performed with the **vpc_pair** policy.



**Note** When Nexus Dashboard Fabric Controller is managing switches, ensure that all configuration changes are initiated from Nexus Dashboard Fabric Controller and avoid making changes directly on the switch.

When the interface config is synced up to the Nexus Dashboard Fabric Controller intent, the switch configs are considered as the reference, that is, at the end of the sync up, the Nexus Dashboard Fabric Controller intent reflects what is present on the switch. If there were any undeployed intent on Nexus Dashboard Fabric Controller for those interfaces before the resync operation, they will be lost.

### Guidelines

- Supported in fabrics using the following templates: Easy_Fabric, External_Fabric, 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.
- The time taken by host port resync depends on the number of switches/interfaces to be synchronized.
- Supported for interfaces that do not have any custom policy (policy template that isn’t shipped with Cisco Nexus Dashboard Fabric Controller) associated with them prior to resync.
- Supported for interfaces where the intent is not exclusively owned by a Cisco Nexus Dashboard Fabric Controller feature and/or application prior to resync.
- Supported on switches that don’t have Interface Groups associated with them.
- Interface mode (switchport to routed, trunk to access, and so on) changes aren’t supported with overlays attached to that interface.

The sync up functionality is supported for the following interface modes and policies:

| Interface Mode | Policies |
|---|---|
| trunk (standalone, po, and vPC PO) | • int_trunk_host<br>• int_port_channel_trunk_host<br>• int_vpc_trunk_host |
| access (standalone, po, and vPC PO) | • int_access_host<br>• int_port_channel_access_host<br>• int_vpc_access_host |

| | |
|---|---|
| dot1q-tunnel | • int_dot1q_tunnel_host<br>• int_port_channel_dot1q_tunnel_host<br>• int_vpc_ dot1q_tunnel_host |
| routed | int_routed_host |
| loopback | int_freeform |
| sub-interface | int_subif |
| FEX (ST, AA) | • int_port_channel_fex<br>• int_port_channel_aa_fex |
| breakout | interface_breakout |
| nve | int_freeform (only in External_Fabric/LAN_Classic) |
| SVI | int_freeform (only in External_Fabric/LAN_Classic) |
| mgmt0 | int_mgmt |

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 Nexus Dashboard Fabric Controller procedures.

## Syncing up Switch Interface Configurations

It is recommended to deploy all switch configurations from NDFC. In some scenarios, it may be necessary to make changes to the switch interface configuration out-of-band. This will cause configuration drift causing switches to be reported Out-of-Sync.

NDFC supports syncing up the out-of-band interface configuration changes back into its intent.

**Guidelines and Limitations**

The following limitations are applicable after Syncing up Switch Interface Configurations to NDFC:

- The port channel membership changes (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** and **LAN_Classic** templates must be updated with the **vpc_pair** policy.
- This feature is supported for easy fabric, external fabric and LAN classic fabric.
- The resync can be performed for a set of switches and repeated as desired.
- The time taken by host port resync depends on the number of switches/interfaces to be synchronized.
- In **Easy_Fabric** fabrics, VXLAN overlay interface attachments are performed automatically based on the allowed VLANs.

### Before you begin

- We recommend taking a fabric backup before attempting the interface resync.
- In **External_Fabric** 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 switch mode must not be **Migration** or **Maintenance**.
- From the **Actions** drop-list, choose **Discovery** > **Rediscover** to ensure that NDFC is aware of any new interfaces and other changes.

### Procedure

**Step 1** Choose **LAN** > **Fabrics** and double-click on a fabric.

The **Fabric Overview** window appears.

**Step 2** Click the **Switches** tab and ensure that switches are present in the fabric and vPC pairings are completed.

**Step 3** Click the **Policies** tab and select one or more switches where the interface intent resync is needed.

**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** From the **Actions** drop-down list, choose **Add Policy**.

The **Create Policy** window appears.

**Step 5** On the **Create Policy** window, choose **host_port_resync** from the **Policy** drop-down list.

**Step 6** Click **Save**.

**Step 7** 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** tab.

**Step 8** After the configuration changes are ready to sync up to NDFC, navigate to the **Switches** tab and select the required switches.

**Step 9** Click **Recalculate & 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.

The time taken by host port resync depends on the number of switches/interfaces to be synchronized.

**Step 10** The **Deploy Configuration** window is displayed if no errors are detected during the resync operation. The interface intent is updated in NDFC.

**Note**
If the External_Fabric or LAN_Classic fabric is in **Monitored Mode**, an error message indicating that the fabric is in the read-only mode is displayed. This error message can be ignored and doesn't mean that the resync process has failed.

Close the **Deploy Configuration** 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.

# Configuration Compliance

The entire intent or expected configuration defined for a given switch is stored in NDFC. 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 configuration 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 NDFC 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 NDFC, CC captures this diff, and indicates that the device is Out-of-Sync. The pending diffs will undo the configurations 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 minutes 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 NDFC REST API Guide*.

To improve ease of use and readability of deployed configurations, CC in NDFC has been enhanced with the following:

- All displayed configurations in NDFC 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.

Top-level configuration commands on the switch that do not have any associated NDFC intent are not checked for compliance by CC. However, CC performs compliance checks, and attempts removal, of the following commands even if there is no NDFC 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** and **Easy_Fabric_eBGP** templates are used. On **External_Fabric** and **LAN_Classic** templates, top-level configuration commands on the switch, including the commands mentioned above, that do not have any associated NDFC intent are not checked for compliance by CC.

We recommend using the NDFC 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.

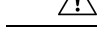
In NDFC, the diffs in default configurations are not highlighted in the **Side-by-side Comparison** window. The auto-generated default configuration that is highlighted in the **Running config** window is not visible in the **Expected config** window.

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 **Recalculate & Deploy**. You need to adhere to the leading space indentations.

You can typically enter configuration snippets in NDFC 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 NDFC 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. Choose **Settings > Server settings**, set the **template.in_use.check** property to **false** and uncheck the **Template In-Use Override** check box and **Save**. This allows the profiles or templates to be editable.
      2. Edit the specific profile or template from the **Operations > Templates > Edit template properties** edit window, and save the updated profile template with the right spacing.
      3. Click **Recalculate & Deploy** to recompute the diffs for the impacted switches.
      4. After the configurations are updated, set the **template.in_use.check** property to **true** and check the **Template In-Use Override** check box and **Save**, as it slows down the performance of the NDFC system, specifically for **Recalculate & Deploy** operations.

If NDFC displays "NA" in the Config Status, the following guidelines apply:

- It is expected when the switch 'Mode' is 'Migration'. This could be due to some of the NDFC work flows. Follow the associated work flow steps to get the switch mode to the 'Normal' state and associated Config Status.
- In all other cases, it may indicate a transient state where NDFC was not able to compute the correct 'Config Status'. Do the following:

- If seen on one switch, then perform switch level **Preview** or **Deploy**.
- If seen on multiple switches, then select those switches and perform **Preview** or **Deploy**.
- If seen at a fabric level, then select all switches and perform **Preview** or **Deploy**.
- R&D is also an option for fabric level but this does a **Config Save** operation as well which could take time in a large fabric.

To confirm that the diffs have been resolved, click **Recalculate & Deploy** after updating the policy to validate the changes.



**Note** NDFC checks only leading spaces, as it implies hierarchy of the command, especially in case of multi-command sequences. NDFC 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 Configuration field.

Create the switch freeform policy.

After deploying this policy successfully to the switch, NDFC persistently reports the diffs.

After clicking the **Side-by-side Comparison** tab, you can see the cause of the diff. The **ip pim rp-address** line has 2 leading spaces, while the running configuration has 0 leading spaces.

To resolve this diff, edit the corresponding Switch Freeform policy so that the spacing is correct.

After you save, you can use the **Push Config** or **Recalculate & Deploy** option to re-compute diffs.

The diffs are now resolved. The **Side-by-side Comparison** tab confirms that the leading spaces are updated.

#### Example 2: Resolving a Leading Space Error in Overlay Configurations

Let us consider an example with a leading space error that is displayed in the **Pending Config** tab.

In the **Side-by-side Comparison** tab, search for diffs line by line to understand context of the deployed configuration.

A matched count of 0 means that it is a special configuration that NDFC has evaluated to push it to the switch.

You can see that the leading spaces are mismatched between running and expected configurations.

Navigate to the respective freeform configs and correct the leading spaces, and save the updated configuration.

Navigate to **Fabric Overview** window for the fabric and click **Recalculate & Deploy**.

In the **Deply Configuration** window, you can see that all the devices are in-sync.

# Configuration Compliance in External Fabrics

With external fabrics, any Nexus switches, Cisco IOS-XE devices, Cisco IOS XR devices, and Arista 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 Nexus Dashboard Fabric Controller, 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 Nexus Dashboard Fabric Controller, 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 Nexus Dashboard Fabric Controller and the switch has additional configuration under the same top-level command, as mentioned earlier, CC will only ensure that the intent defined in Nexus Dashboard Fabric Controller is present on the switch. When this user defined intent on Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller and deployed to the switch.
2. Additional configuration exists under **router bgp** in **Running config** that does not exist in user-defined Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller.
3. The **Pending Config** and the **Side-by-side Comparison** when the intent that was pushed earlier via Nexus Dashboard Fabric Controller is deleted from Nexus Dashboard Fabric Controller by deleting the switch_freeform policy that was created in the Step 1.
4. A **switch_freeform** policy with the top-level **router bgp** command needs to be created. This enables CC to generate the configuration needed to remove only the desired sub-config which was pushed from Nexus Dashboard Fabric Controller earlier.
5. The removed configuration is only the subset of the configuration that was pushed earlier from Nexus Dashboard Fabric Controller.

   For interfaces on the switch in the external fabric, Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller. 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:

### Procedure

**Step 1** Select **LAN > Fabrics**.

Double click on the fabric name to view **Fabric Overview** screen.

**Step 2** On the Switches tab, double click on the switch name to view **Switch Overview** screen.

On the Policies tab, all the policies applied on the switch within the chosen fabric are listed.

**Step 3** On the Policies tab, from the **Actions** drop-down list, select **Add Policy**.

**Step 4** Add a Policy Template Instances (PTIs) with the required configuration CLIs using the **switch_freeform** template and click **Save**.

**Step 5** Select the created policy and select **Push Config** from the **Actions** drop-down list to deploy the configuration to the switch(es).

## Resolving Diffs for Case Insensitive Commands

By default, all diffs generated in NDFC 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 are captured in the **compliance_case_insensitive_clis.txt** template that can be found under **Operations** > **Templates**.

From Cisco NDFC Release 12.0.1a, **compliance_case_insensitive_clis.txt** file, along with **compliance_strict_cc_exclude_clis.txt** and **compliance_ipv6_clis.txt** files are now part of the shipped templates. You can find all the templates under **Operations** > **Templates**. Modification of templates can be done after disabling **Template In-Use Override**.

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 NDFC and the Running Configuration, you can configure NDFC to ignore these case differences as follows:

1. Navigate to **Settings** > **Server Settings** > **LAN-Fabric**, uncheck **Template In-Use Override** and then click **Save**.
2. Navigate to **Operations** > **Templates** and search for **compliance_case_insensitive_clis.txt** file.
3. Check **compliance_case_insensitive_clis.txt** and choose **Actions** > **Edit template content**.

   An example of the entries in the **compliance_case_insensitive_clis.txt** file is displayed in the following figure.
4. Remove the entries highlighted in the figure and click **Finish**.

```
##template variables
##
##template content
"^(no |)interface\s+Port(.)"
"^(no |)interface\s+Loo(.)"
"^(no |)interface\s+Eth(.)"
"^update-source\s+Loo(.)"
"^vrf\s+"
"^destination\s(.+)\suse-vrf\s(.+)"
"^hardware profile portmode\s+"
"^(?!.*(ospf|isis|bgp)(?:|$))(.*)route-map\s+(.)"
"^(.*)neighbor-policy(.)"
"(no |)encapsulation\s+(.)"
"(.*)alert-group\s+(.)"
"^streetaddress\s+(.)"
"^transport email\s+(.)"
"(no |)action\s+(.)"
"(no |)\d+\s+remark.*"
"(no |)\d+\s+permit.*"
"(no |)\d+\s+deny.*"
"(no |)ip(v6|)\s+access-list.*"
"(no |)ip\s+access-group.*"
"(no |)ipv6\s+traffic-filter.*"
"^mac-address\s+([A-Fa-f0-9]{4}[.]?){3}"
"\s*ip\s+dhcp\s+relay\s+address\s+(.)"
##

```

5. 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.
6. Navigate to **Settings** > **Server Settings** > **LAN-Fabric**, check **Template In-Use Override** and then click **Save**.

   This enables NDFC to treat the documented configuration patterns as case insensitive while performing comparisons.
7. Click **Recalculate & Deploy** for fabrics to view the updated comparison outputs.

## Resolving Configuration Compliance After Importing Switches

After importing switches in Cisco NDFC, 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 configuration profile:

```
interface mgmt0
  description SRC=SDS-LB-LF111-mgmt0,DST=SDS-LB-SW001-Fa0/5
```

Navigate to Interface Manager and click the **Edit** icon after selecting the mgmt0 interface. Remove the extra space in the description.

## Strict Configuration Compliance

Strict configuration compliance checks for diff between the switch configuration and the associated intent and generates **no** commands for the configurations that present on the switch but are not present in the associated intent. When you click **Recalculate and Deploy**, switch configurations that are not present on the associated intent are removed. You can enable this feature by choosing the **Enable Strict Config Compliance** check box under the **Advanced** tab in the **Create Fabric** or **Edit Fabric** window. By default, this feature is disabled.

The strict configuration compliance feature is supported on the Easy Fabric templates - **Easy_Fabric** 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 maintained in **Operations** > **Templates**, **compliance_strict_cc_exclude_clis.txt** template.

### 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 configuration compliance feature is enabled, the **no** form of the **feature telnet** command appears under **Pending Config** in the **Preview Config** window.

Click the **Side-by-side Comparison** tab to display the differences between the running configuration and the expected configuration. 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 NDFC state when there is a large scale out-of-band change, or if configuration changes do not register in the NDFC 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. 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 NDFC.

Now, close the **Preview Config** window and click **Recalculate and Deploy**. The strict configuration compliance feature ensures that the running configuration 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.

You can right-click on a switch in the **Fabric Overview** 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 NDFC to make the intended configuration on NDFC and the switch configurations identical. The switches are then in In-Sync status. For more information on how to add custom freeform configurations on NDFC, refer Enabling Freeform Configurations on Fabric Switches , on page 95.

# Enabling Freeform Configurations on Fabric Switches

In Nexus Dashboard Fabric Controller, 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.
4. On a specific interface on a switch.

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. Choose **LAN** > **Fabrics** > **Fabrics**.
2. Select the Fabric, and select **Edit Fabric** from **Actions** drop-down list.

   (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 98.

4. Click **Save**. The fabric topology screen comes up.
5. Click **Deploy Config** from the **Actions** drop-down list to save and deploy configurations.

   Configuration Compliance functionality ensures 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 flags 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 **Deploy Config** 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 **Deploy Config** in the topology screen to complete the deployment process.

To bring back the switch 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. Choose **LAN > Fabrics > Fabrics**.
2. Select the Fabric, and select **Edit Fabric** from **Actions** drop-down list.
3. Click **Policies** tab. From the **Actions** drop-down list, choose **Add Policy**.

   The **Create Policy** 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.

4. 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. In the **Description** field, provide a description for the policy.
6. From the **Template Name** field, select **freeform_policy**.
7. 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 98.

8. Click **Save**.

   After the policy is saved, it gets added to the intended configurations for that switch.

9. From the Fabric Overview window, click the **Switches** tab and choose the required switches.

10. On the **Switches** tab, click **Actions** drop-down list and choose **Deploy**.

**Pointers for freeform_policy Policy Configuration:**

- You can create multiple instances of the policy.
- For a vPC switch pair, create consistent **freeform_policy** policies on both the vPC switches.
- When you edit a **freeform_policy** 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 configuration 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:

```
ip prefix-list prefix-list-name1 seq 5 permit 20.2.0.1/32
ip prefix-list prefix-list-name1 seq 6 permit 20.2.0.2/32
ip prefix-list prefix-list-name2 seq 5 permit 192.168.100.0/24
```

### 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 **freeform_policy** 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.

### 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 Nexus Dashboard Fabric Controller 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 **Recalculate Config**, click the **Pending Config** column. The **Side-by-Side Comparison** to view 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. Choose **LAN > Fabrics > Fabrics**.
2. Select the Fabric, and select **Edit Fabric** from **Actions** drop-down list.
3. Click **VRFs** tab. From the **Actions** drop-down list, select **Create**.

   The **Create VRF** screen comes up.
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 shows up.
5. Under the CLI Freeform column, select the button labeled **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 **Deploy 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 Config** to push the configuration to the switch.

The same procedure can be used to define a per Network per Switch configuration.

# MACsec Support in Easy Fabric and eBGP Fabric

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).

### 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 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 Nexus Dashboard Fabric Controller:

## Enabling MACsec

### Procedure

**Step 1** Navigate to **LAN** > **Fabrics**.

**Step 2** Click **Actions** > **Create** to create a new fabric or click **Actions** > **Edit Fabric** on an existing Easy or eBGP fabric.

**Step 3** Click the **Advanced** tab and specify the MACsec details.

**Enable MACsec** – Select the check box to enable MACsec for the fabric.

**MACsec Primary Key String** – Specify a Cisco Type 7 encrypted octet string that is used for establishing the primary MACsec session. For AES_256_CMAC, the key string length must be 130 and for AES_128_CMAC, the key string length must be 66. If these values are not specified correctly, an error is displayed when you save the fabric.

**Note**
The default key lifetime is infinite.

**MACsec Primary Cryptographic Algorithm** – Choose the cryptographic algorithm used for the primary key string. It can be **AES_128_CMAC** or **AES_256_CMAC**. The default value is **AES_128_CMAC**.

You can configure a fallback key on the device to initiate a backup session if the primary session fails.

**MACsec Fallback Key String** – Specify a Cisco Type 7 encrypted octet string that is used for establishing a fallback MACsec session. For AES_256_CMAC, the key string length must be 130 and for AES_128_CMAC, the key string length must be 66. If these values are not specified correctly, an error is displayed when you save the fabric.

**MACsec Fallback Cryptographic Algorithm** – Choose the cryptographic algorithm used for the fallback key string. It can be AES_128_CMAC or AES_256_CMAC. The default value is AES_128_CMAC.

**MACsec Cipher Suite** – Choose one of the following MACsec cipher suites for the MACsec policy:

- GCM-AES-128
- GCM-AES-256

- GCM-AES-XPN-128
- GCM-AES-XPN-256

The default value is **GCM-AES-XPN-256**.

**Note**
The MACsec configuration is not deployed on the switches after the fabric deployment is complete. You need to enable MACsec on intra-fabric links to deploy the MACsec configuration on the switch.

**MACsec Status Report Timer** – Specifies MACsec operational status periodic report timer in minutes.

**Step 4** Click a fabric to view the **Summary** in the side kick. Click the side kick to expand. Click **Links** tab.

**Step 5** Choose an intra-fabric link on which you want to enable MACsec and click **Actions > Edit**.

**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** From the Fabric Actions drop-down list, select **Deploy Config** 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**. From the Fabric Actions drop-down list, select **Deploy Config** to disable MACsec configuration. 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 **Actions > Recalculate Config** from the **Fabric Actions** drop-down list.

# Create Easy_Fabric for Cisco Catalyst 9000 Series Switches

You can add Cisco Catalyst 9000 Series Switches to an easy fabric using the Easy_Fabric_IOS_XE fabric template. You can add only Cisco Catalyst 9000 IOS XE switches to this fabric. This fabric supports OSPF as underlay protocol and BGP EVPN as the overlay protocol. Using this fabric template allows Nexus Dashboard Fabric Controller to manage all the configurations of a VXLAN EVPN Fabric composed of Cisco Catalyst 9000 IOS-XE switches. Backing up and restoring this fabric is the same as the **Easy_Fabric**.

**Guidelines**

- EVPN VXLAN Distributed Anycast Gateway is supported when each SVI is configured with the same Anycast Gateway MAC.
- StackWise Virtual switch is supported.
- Brownfield is not supported.
- Upgrade from earlier versions is not supported (However, it is a preview feature in 11.5).
- IPv6 Underlay, VXLAN Multi-site, Anycast RP, and TRM is not supported.
- ISIS, ingress replication, unnumbered intra-fabric link, 4 bytes BGP ASN, and Zero-Touch Provisioning (ZTP) is not supported.



**Note** For information about configuration compliance, see Configuration Compliance in External Fabrics, on page 91.

### Creating Easy Fabric for Cisco Catalyst 9000 Series Switches

**UI Navigation**: Choose **LAN > Fabrics**.

Perform the following steps to create an easy fabric for Cisco Catalyst 9000 Series Switches:

1. Choose **Create Fabric** from the **Actions** drop-down list.
2. Enter a fabric name and click **Choose Template**.

   The **Select Fabric Template** dialog appears.
3. Choose the **Easy_Fabric_IOS_XE** fabric template and click **Select**.
4. Fill in all the required fields and click **Save**.



**Note** BGP ASN is the only mandatory field.

## Adding Cisco Catalyst 9000 Series Switches to IOS-XE Easy Fabrics

Cisco Catalyst 9000 series switches are discovered using SNMP. Hence, before adding them to the fabric, configuring the Cisco Catalyst 9000 series switches includes configuring SNMP views, groups, and users. For more information, see the Configuring IOS-XE Devices for Discovery section.

For StackWise Virtual switches, configure the StackWise Virtual-related configuration before adding them to the fabric.

### UI Navigation

Choose any one of the following navigation paths to add switch(es) in the **Add Switches** window.

- Choose **LAN > Fabrics**. Choose a fabric that uses the **Easy_Fabric_IOS_XE** fabric template from the list, click **Actions**, and choose **Add Switches**.
- Choose **LAN > Fabrics**. Choose a fabric that uses the **Easy_Fabric_IOS_XE** fabric template from the list. Click the **Switches** tab. Click **Actions** and choose **Add Switches**.

- Choose **LAN > Switches**. Click **Actions** and choose **Add Switches**. Click **Choose Fabric**, choose the IOS-XE VXLAN fabric, and click **Select**.

### Before you begin

Set the default credentials for the device in the **LAN Credentials Management** window if the default credentials are not set. To navigate to the **LAN Credentials Management** window from the Cisco Nexus Dashboard Fabric Controller Web UI, choose **Settings** > **LAN Credentials Management**.

### Procedure

**Step 1** Enter values for the following fields:

| Field | Description |
|---|---|
| Seed IP | Enter the IP address of the switch.<br>You can import more than one switch by providing the IP address range. For example: 10.10.10.40-60<br>The switches must be properly cabled and reachable to the Cisco Nexus Dashboard Fabric Controller server and the switch status must be manageable. |
| Authentication Protocol | Choose the authentication protocol from the drop-down list. |
| Username | Enter the username of the switch(es). |
| Password | Enter the password of the switch(es). |

**Note**
You can change the Discover and LAN credentials only after discovering the switch.

**Step 2** Click **Discover Switches**.

The switch details are populated.

Cisco Nexus Dashboard Fabric Controller supports the import of Cisco Catalyst 9500 Switches running in StackWise Virtual. The StackWise Virtual configuration to form a pair of Cisco Catalyst 9500 Switches into a virtual switch has to be in place before the import. For more information on how to configure StackWise Virtual, see the Configuring Cisco StackWise Virtual chapter in the *High Availability Configuration Guide (Catalyst 9500 Switches)* for the required release.

**Step 3** Check the check boxes next to the switches you want to import.

You can import only switches with the **manageable** status.

**Step 4** Click **Add Switches**.

The switch discovery process is initiated and the discovery status is updated under the **Discovery Status** column in the **Switches** tab.

**Step 5** (Optional) View the details of the device.

After the discovery of the device, the discovery status changes to **ok** in green.

### What to do next

1. Set the appropriate role. The supported roles are:
   - Leaf
   - Spine
   - Border

   To set the role, choose a switch and click **Actions**. Choose **Set role**. Choose a role and click **Select**.



**Note** After discovering the switch(es), Nexus Dashboard Fabric Controller usually assigns **Leaf** as the default role.

2. Recalculate the configurations and deploy the configurations to the switches.

## Recalculating and Deploying Configurations

To recalculate and deploy the configurations to the switch(es) in the IOS-XE easy fabric, perform the following steps to recalculate configurations:

### Before you begin

Set the role of the switch(es) in the IOS-XE easy fabric.

### Procedure

**Step 1** Click **Actions** from **Fabric Overview**.

**Step 2** Choose **Recalculate Config**.

Recalculation of configurations starts on the switch(es).

## Creating DCI Links for Cisco Catalyst Switches in IOS-XE Easy Fabrics

You can create VRF-Lite IFC between a Cisco Catalyst 9000 Series Switch with border role in IOS-XE easy fabrics, and another switch in a different fabric. The other switch can be a Nexus switch in External Fabric, LAN Classic fabric, or Easy Fabric. It can also be a Catalyst 9000 switch in External Fabric or IOS-XE Easy Fabric. The link can be created only from IOS-XE Easy Fabric.

For more information, see Links, on page 188 and Templates, on page 413.

**Note** When creating DCI links for IOS-XE Easy Fabric, auto-deploy is supported only if the destination device is a Nexus switch.

To create links for IOS-XE Easy Fabric, perform the following procedure:

1. Navigate to the **Links** tab in the fabric overview.

   The list of previously created links is displayed. The list contains intra-fabric links, which are between switches within a fabric, and inter-fabric links, which are between border switches in this fabric and switches in other fabrics.

   The inter-fabric links also support edge router switches in the External Fabric, apart from BGW and Border Leaf/Spine.

2. Click **Actions** and choose **Create**.

   The **Create Link** window appears. By default, the **Intra-Fabric** option is chosen as the link type.

3. From the **Link Type** drop-down box, choose **Inter-Fabric** . The fields change correspondingly.

4. Choose **VRF_LITE** as the link sub-type, ext_fabric_setup template for VRF_LITE IFC, and IOS-XE fabric as the source fabric.

   Link Template: The link template is populated.

   The templates are autopopulated with corresponding pre-packaged default templates that are based on your selection. The template to use for VRF_LITE IFC is ext_fabric_setup.

**Note** You can add, edit, or delete only the ext_routed_fabric template. For more information, see Templates.

5. Choose the IOS-XE fabric as the source fabric from the Source Fabric drop-down list.
6. Choose a destination fabric from the Destination Fabric drop-down list.
7. Choose the source device and Ethernet interface that connects to the destination device.
8. Choose the destination device and Ethernet interface that connects to the source device.
9. Enter values in other fields accordingly.
10. Click **Save**.

**Note** Instead of the create action, you can also use the **Edit** action to create VRF-Lite IFC(s) using the existing inter fabric link(s). Choose the **VRF_Lite** link subtype. By default, if you select **Edit**, then the data for the fields Link-Type, Source Fabric, Destination Fabric, Source Device, Destination Device, Source Interface and Destination Interface are auto-populated in the **Edit Link** window.

Choose **VRF_LITE** as the link sub-type, ext_fabric_setup template for VRF_LITE IFC, and IOS-XE fabric as the source fabric.

To complete the procedure, repeat step 4 to step 10 mentioned above.

# Creating VRFs for Cisco Catalyst 9000 Series Switches in IOS-XE Easy Fabrics

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **VRFs** > **VRFs**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **VRFs** > **VRFs**.

You can create VRFs for IOS-XE easy fabrics.

To create VRF from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Click **Actions** and choose **Create**.

   The **Create VRF** window appears.

2. Enter the required details in the mandatory fields. Some of the fields have default values.

   The fields in this window are:

   **VRF Name** - Specifies a VRF name automatically or allows you to enter a name for Virtual Routing and Forwarding (VRF). The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).

   **VRF ID** - Specifies the ID for the VRF or allows you to enter an ID for the VRF.

   **VLAN ID** - Specifies the corresponding tenant VLAN ID for the network or allows you to enter an ID for the VLAN. If you want to propose a new VLAN for the network, click **Propose Vlan**.

   **VRF Template** - A universal template is autopopulated. This is only applicable for leaf switches. The default template for IOS_XE Easy Fabric is the **IOS_XE_VRF** template.

   **VRF Extension Template** - A universal extension template is autopopulated. This allows you to extend this network to another fabric. The default template for IOS_XE Easy Fabric is the **IOS_XE_VRF** template.

   The VRF profile section contains the **General Parameters** and **Advanced** tabs.

3. The fields on the **General** tab are:

   **VRF Description** - Enter the a description for the VRF.

   **VRF Intf Description** - Specifies the description for the VRF interface.

4. Click the **Advanced** tab to optionally specify the advanced profile settings. The fields on the **Advanced** tab are:

   **Redistribute Direct Route Map** - Specifies the redistribute direct route map name.

   **Max BGP Paths** - Specifies the maximum BGP paths. The valid value range is between 1 and 64.

   **Max iBGP Paths** - Specifies the maximum iBGP paths. The valid value range is between 1 and 64.

   **Advertise Host Routes** - Enable this check box to control advertisement of /32 and /128 routes to Edge routers.

   **Advertise Default Route** - Enable this check box to control advertisement of default route internally.

   **Config Static 0/0 Route** - Enable this check box to control configuration of static default route.

5. Click **Create** to create the VRF or click **Cancel** to discard the VRF.

   A message appears indicating that the VRF is created.

The new VRF appears on the **VRFs** horizontal tab. The status is **NA** as the VRF is created but not yet deployed. Now that the VRF is created, you can create and deploy networks on the devices in the fabric.

**What to do next**

Attach the VRF.

Create a loopback interface selecting the VRF_LITE extension.

For more information about attaching and detaching VRFs, see VRF Attachments, on page 207.

## Attaching VRFs on Cisco Catalyst 9000 Series Switches in IOS-XE Easy Fabrics

To attach the VRFs on the Cisco Catalyst 9000 Series Switches in the IOS-XE easy fabric, see VRF Attachments, on page 207.

**Note** Choose the VRF corresponding to the CAT9000 series switch by checking the check box next to it.

**Note** Similarly, you can create a loopback interface, and select VRF_LITE extension.

**What to do next**

Deploy the configurations as follows:

1. Click **Actions** in **Fabric Overview**.
2. Choose **Deploy config to switches**.
3. Click **Deploy** after the configuration preview is complete.
4. Click **Close** after the deployment is complete.

## Creating and Deploying Networks in IOS-XE Easy Fabrics

The next step is to create and deploy networks in IOS-XE Easy Fabrics.

**Note**

- The Network Template and Network Extension template uses the default IOS_XE_Network template that was created for the IOS-XE easy fabric.

**UI Navigation**

The following options are applicable only for switch fabrics, easy fabrics, and MSD fabrics:

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Networks**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Networks**.

### Creating Networks for IOS-XE Easy Fabrics

To create network for IOX-XE easy fabric from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. On the **Networks** horizontal tab, click **Actions** and choose **Create**.

   The **Create Network** window appears.

2. Enter the required details in the mandatory fields.

   The fields in this window 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 (-).

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

   **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 **Create VRF**. The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).

   **VLAN ID** - Specifies the corresponding tenant VLAN ID for the network. If you want to propose a new VLAN for the network, click **Propose VLAN**.

   **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 VRF Lite extension is supported. The template is applicable for border leaf switches.

   **Generate Multicast IP** - If you want to generate a new multicast group address and override the default value, click **Generate Multicast IP**.

   The network profile section contains the **General** and **Advanced** tabs.

3. The fields on the **General** tab are:



**Note** If the network is a non Layer 2 network, then it is mandatory to provide the gateway IP address.

**IPv4 Gateway/NetMask** - Specifies the IPv4 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. The anycast gateway IP address is the same for MyNetwork_30000 on all switches of the fabric that have the presence of the network.



**Note** If the same IP address is configured in the IPv4 Gateway and IPv4 Secondary GW1 or GW2 fields of the network template, Nexus Dashboard Fabric Controller 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 List** - Specifies the IPv6 address with subnet.

**Vlan Name** - Enter the VLAN name.

**Vlan Interface Description** - Specifies the description for the interface. This interface is a switch virtual interface (SVI).

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

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

4. Click the **Advanced** tab to optionally specify the advanced profile settings. The fields on the **Advanced** tab are:

   **Multicast Group Address** - The multicast IP address for the network is autopopulated.

   Multicast group address is a per fabric instance variable and remains the same for all networks by default. If a new multicast group address is required for this network, 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 VRF** - Enter the DHCP server VRF ID.

   **DHCPv4 Server 2** - Enter the DHCP relay IP address of the next DHCP server.

   **DHCPv4 Server2 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.

   **Enable L3 Gateway on Border** - Select the check box to enable a Layer 3 gateway on the border switches.

5. Click **Create**.

   A message appears indicating that the network is created.

   The new network appears on the **Networks** page that comes up.

   The Status is **NA** since the network is created but not yet deployed on the switches. Now that the network is created, you can create more networks if needed and deploy the networks on the devices in the fabric.

### Deploying Networks in IOS-XE Easy Fabrics

You can deploy networks in IOS-XE easy fabrics as follows:

- The network configurations can also be deployed in the **Fabric Overview** window as follows:
  1. Click **Actions** in the fabric overview.
  2. Choose **Deploy config to switches**.
  3. Click **Deploy** after the configuration preview is complete.
  4. Click **Close** after the deployment is complete
- To deploy the network in the IOS-XE easy fabric, see Network Attachments, on page 216.

# External Fabrics

You can add switches to the external fabric. Generic pointers:

NDFC will not generate "no router bgp". If you want to change it, go to the switch and do a “no feature bgp” followed by a re-sync, if you don't have anything and want to update the ASN.

- The external fabric is a monitor-only or managed mode fabric.
- From Cisco Nexus Dashboard Fabric Controller Release 12.0.1, Cisco IOS-XR family devices Cisco ASR 9000 Series Aggregation Services Routers and Cisco Network Convergence System (NCS) 5500 Series are supported in external fabric in managed mode and monitor mode. NDFC will generate and push configurations to these switches, and configuration compliance will also be enabled for these platforms.
- 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-Nexus Dashboard Fabric Controller 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.
- 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

- You can configure a Cisco ASR 1000 Series routers and Cisco Catalyst 9000 Series switches as edge routers, set up a VRF-lite IFC and connect it as a border device with an easy fabric.
- Before a VDC reload, discover Admin VDC in the fabric. Otherwise, the reload operation does not occur.
- You can connect a Cisco data center to a public cloud using Cisco CSR 1000v. See the *Connecting Cisco Data Center and a Public Cloud* chapter for a use case.
- In an external fabric, when you add the **switch_user** policy and provide the username and password, the password must be an encrypted string that is displayed in the **show run** command.

  For example:

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

  In this case, the entered password should be **5$5$I4sapkBh$S7B7UcPH/iVTihLKH5sgldBeS3O2X1StQsvv3cmbYd1**.
- For the Cisco Network Insights for Resources (NIR) Release 2.1 and later, and flow telemetry, **feature lldp** command is one of the required configuration.

  Cisco Nexus Dashboard Fabric Controller 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
- Backup/restore is only supported for Nexus devices on external fabrics.



**Note** Before you do fabric or switch restore, ensure that the target device is supported. If the target device is not supported, then per switch restore will be blocked, and the same will be shown as not supported during fabric-wide restore.

### 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. On **Topology**, click within the MSD-Parent-Fabric. From **Actions** drop-down list, select **Move Fabrics**.

   The Move Fabric screen comes up. It contains a list of fabrics. The external fabric is displayed as a standalone fabric.
2. 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.

### 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.

## Creating an External Fabric

To create an external fabric using Cisco Fabric Controller Web UI, perform the following steps:

### Procedure

**Step 1** Choose **LAN > Fabrics > Fabrics**.

**Step 2** From the **Actions** drop-down list, select **Create Fabric**.

**Step 3** Enter a unique name for the fabric and click **Choose Template**.

**Step 4** From the drop-down list, select **External_Fabric** template.

The fields in this screen are:

**BGP AS #** – Enter the BGP AS number.

**Fabric Monitor Mode** – Clear the check box if you want Nexus Dashboard Fabric Controller to manage the fabric. Keep the check box selected to enable a monitor only external fabric.

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 **Deploy Config**, 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.

**Enable Performance Monitoring** – Check this check box to enable performance monitoring on NX-OS switches only.

Ensure that you do not clear interface counters from the Command Line Interface of the switches. Clearing interface counters can cause the Performance Monitor to display incorrect data for traffic utilization. If you must clear the counters and the switch has both `clear counters` and `clear counters snmp` commands (not all switches have the `clear counters snmp` command), ensure that you run both the main and the SNMP commands simultaneously. For example, you must run the `clear counters interface ethernet` *`slot/port`* command followed by the `clear counters interface ethernet` *`slot/port`* `snmp` command. This can lead to a one time spike.

**Step 5** Enter values in the fields under the **Advanced** tab.

**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 MPLS SR and LDP Handoff, on page 619 chapter in External/WAN Layer 3 Connectivity for VXLAN BGP EVPN Fabrics.

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

**Enable AAA IP Authorization** – Enables AAA IP authorization, after IP Authorization is enabled on the AAA Server

**Enable** Nexus Dashboard Fabric Controller **as Trap Host** – Select this check box to enable Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller, 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 Nexus Dashboard Fabric Controller to import and manage of switches with inband connectivity (reachable over switch loopback, or routed interface, 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 Nexus Dashboard Fabric Controller to the switches over the Nexus Dashboard data interface, also known as inband interface. For this purpose, static routes may be needed on the Nexus Dashboard Fabric Controller, that in turn can be configured from **Administration > Customization > Network Preferences**. 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. Nexus Dashboard Fabric Controller has a precheck that validates that the Inband managed switch IPs are reachable over the Nexus Dashboard data interface. After completing the precheck, Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller. For more information, see Inband Management in External Fabrics and LAN Classic Fabrics, on page 152.

**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 Nexus Dashboard Fabric Controller are typically bound to the eth1 or out-of-band interface. In scenarios, where Nexus Dashboard Fabric Controller 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. You can also edit **PTP Source Loopback Id** and **PTP Domain Id** fields. For more information, see Precision Time Protocol for External Fabrics and LAN Classic Fabrics, on page 151.

**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 0–1023. 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 Nexus Dashboard Fabric Controller. 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 0–127.

**Fabric Freeform** – You can apply configurations globally across all the devices that are 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 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.

**AAA Freeform Config** – You can apply AAA configurations globally across all devices that are discovered in the external fabric using this freeform field.

**Step 6** Fill up the **Resources** tab as explained in the following.

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

**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.

**Step 7** Fill up the **Configuration Backup** tab as shown below.

The fields on this tab are:

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

You can enable an hourly backup for fresh fabric configurations and the intent as well. If there is a configuration push in the previous hour, Nexus Dashboard Fabric Controller takes a backup. In case of the external fabric, the entire configuration on the switch is not converted to intent on Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller 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 that 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 Fabric** in the **Actions** pane.

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

**Step 8** Click the **Bootstrap** tab.

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

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

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

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

**Note**
Cisco Nexus Dashboard Fabric Controller 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, Nexus Dashboard Fabric Controller 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 range is 8-30.

*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 from 112 through 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 that is 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 Enabling Freeform Configurations on Fabric Switches , on page 95.

**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

**Step 9** Click the **Flow Monitor** tab. The fields on this tab are as follows.

**Enable NetFlow** – Check this check box to enable NetFlow on VTEPs for this Fabric. By default, NetFlow is disabled. On Enable, NetFlow configuration will be applied to all VTEPS that support NetFlow.

**Note:** When NetFlow is enabled on the fabric, you can choose not to have NetFlow on a particular switch by having a dummy no_netflow PTI.

If NetFlow is not enabled at the fabric level, an error message is generated when you enable NetFlow at the interface, network, or VRF level. For information about NetFlow support for Cisco NDFC, see Netflow Support, on page 149.

In the **NetFlow Exporter** area, click **Actions > Add** to add one or more NetFlow exporters. This exporter is the receiver of the NetFlow data. The fields on this screen are:

- **Exporter Name** – Specifies the name of the exporter.
- **IP** – Specifies the IP address of the exporter.
- **VRF** – Specifies the VRF over which the exporter is routed.
- **Source Interface** – Enter the source interface name.
- **UDP Port** – Specifies the UDP port over which the NetFlow data is exported.

Click **Save** to configure the exporter. Click **Cancel** to discard. You can also choose an existing exporter and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

In the **NetFlow Record** area, click **Actions > Add** to add one or more NetFlow records. The fields on this screen are:

- **Record Name** – Specifies the name of the record.
- **Record Template** – Specifies the template for the record. Enter one of the record templates names. In Release 12.0.2, the following two record templates are available for use. You can create custom NetFlow record templates. Custom record templates that are saved in the template library are available for use here.
  - **netflow_ipv4_record** – to use the IPv4 record template.
  - **netflow_l2_record** – to use the Layer 2 record template.
- **Is Layer 2 Record** – Check this check box if the record is for Layer 2 NetFlow.

Click **Save** to configure the report. Click **Cancel** to discard. You can also choose an existing record and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

In the **NetFlow Monitor** area, click **Actions > Add** to add one or more NetFlow monitors. The fields on this screen are:

- **Monitor Name** – Specifies the name of the monitor.
- **Record Name** – Specifies the name of the record for the monitor.
- **Exporter1 Name** – Specifies the name of the exporter for the NetFlow monitor.
- **Exporter2 Name** – (optional) Specifies the name of the secondary exporter for the NetFlow monitor.

The record name and exporters referred to in each NetFlow monitor must be defined in **Netflow Record** and **Netflow Exporter**.

Click **Save** to configure the monitor. Click **Cancel** to discard. You can also choose an existing monitor and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

**Step 10** Click **Save**.

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

After creating the external fabric, add switches to it.

## Adding Switches to the External Fabric

Switches in each fabric are unique, and hence, each switch can only be added to one fabric. To add switches to the external fabric, perform he following steps:

### Procedure

**Step 1** Choose **LAN > Switches**. From the Actions drop-down list, select **Add Switches**

You can also add switches to a Fabric from **LAN > Fabrics**. Select a fabric and view the **Summary**. On the **Switches** tab, from the **Actions** drop-down list, select **Add switches** to add switches to the selected Fabric.

From Topology, right click on the Fabric and select **Add Switches**.

**Step 2** Select **Discover** to discover new switches. Select **Move Neighbor Switches** to add existing switches to the Fabric.

**Step 3** If you select **Discover** option, perform the following steps:

a) Enter the IP address (Seed IP) of the switch.
b) In the **Authentication Protocol** field, from the drop-down list, select the appropriate protocol to add switches to the Fabric.
c) Choose the device type from the **Device Type** drop-down list.

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

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

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.

d) Enter the administrator username and password of the switch.
e) Click **Discovery Switches** 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.

Select the check boxes next to the concerned switches and click **Add Switches** into fabric.

You can discover multiple switches at the same time. The switches must be properly cabled and connected to the Nexus Dashboard Fabric Controller server and the switch status must be manageable.

The switch discovery process is initiated. The **Progress** column displays the progress. After Nexus Dashboard Fabric Controller discovers the switch, click **Close** to revert to the previous screen.

**Step 4** If you select **Move Neighbor Switches** option, select the switch and click **Move Switch**.

The selected switch is moved to the External Fabric.

## Switch Settings for External Fabrics

External Fabric Switch Settings vary from the VXLAN fabric switch settings. Double-click on the switch to view the Switch Overview screen to edit/modify options.

The options are:

**Set Role** – By default, no role is assigned to an external fabric switch. You can assign desired role to the fabric. 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 **Deploy Config**.

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

**Change Modes** – Allows you to modify the mode of switch from Active to Operational.

**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.

**Recalculate 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.

Click **Deploy** from the Actions drop-down list. The template and interface configurations form the configuration provisioning on the switches.

When you click **Deploy**, the **Deploy Configuration** screen comes up.

Click **Config** at the bottom part of the screen to initiate pending configuration onto the switch. The **Deploy Progress** screen displays the progress and the status of configuration deployment.

Click **Close** after the 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, Nexus Dashboard Fabric Controller discovery continues, but the switch status shows a warning for the SSH error.

## 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 Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller UI, select the External Fabric. Choose **Edit Fabric** from the **Actions** drop-down list.

The **Edit Fabric** screen is displayed.

**Step 4** Click the **Bootstrap** tab and update the DHCP information.

**Step 5** Click **Save** at the bottom right part of the **Edit Fabric** screen to save the settings.

**Step 6** Double click on the Fabric to view the **Fabric Overview**.

**Step 7** On **Switches** tab, from the **Actions** drop-down list, select **Add Switches**.

**Step 8** Click the **POAP** tab.

In an earlier step, the reload command was executed on the switch. When the switch restarts to reboot, Nexus Dashboard Fabric Controller retrieves the serial number, model number, and version from the switch and

displays them on the Inventory Management along screen. Also, an option to add the management IP address, hostname, and password are made available. If the switch information is not retrieved, refresh the screen using the Refresh icon at the top right part of the screen.

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

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

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

You can provision devices in advance.

**Step 9** 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 10** (Optional) Use discovery credentials for discovering switches.

a) Click the **Add Discovery Credentials** icon to enter the discovery credentials for switches.
b) In the **Discovery Credentials** window, enter the discovery credentials such as discovery username and password.

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

If the discovery credentials are not provided, Nexus Dashboard Fabric Controller 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 11** Click **Bootstrap** at the top right part of the screen.

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

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

**Step 12** Monitor and check the switch for POAP completion.

**Step 13** Click **Deploy Config** from the **Actions** drop-down list on the **Fabric Overview** 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 Nexus Dashboard Fabric Controller, 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 14** After the pending configurations are deployed, the **Progress** column displays 100% for all switches.

**Step 15** On the Topology screen, click **Refresh Topology** icon to view the update.

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

The switch and the link are discovered in Nexus Dashboard Fabric Controller. 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 16** Right-click and select History to view the deployed configurations.

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

**Step 17** On the Nexus Dashboard Fabric Controller 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 **LAN > 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

From Cisco Nexus Dashboard Fabric Controller Release 12.0.1a, you can add Cisco IOS-XR devices to external fabrics in managed mode as well. You can manage the following Cisco IOS-XR devices in external fabrics:

- Cisco ASR 9000 Series Routers

- Cisco NCS 5500 Series Routers, IOS XR Release 6.5.3
- Cisco 8000 Series Routers

You can discover non-Nexus devices in an external fabric and perform the configuration compliance of these devices as well. For more information, see the Configuration Compliance in External Fabrics, on page 91 section.

Refer the *Cisco Nexus Dashboard Fabric Controller Compatibility Matrix* to see the non-Nexus devices supported by Cisco Nexus Dashboard Fabric Controller.

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 Nexus Dashboard Fabric Controller can only access the basic device information like system name, serial number, model, version, interfaces, up time, and so on. Cisco Nexus Dashboard Fabric Controller 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.

You can add IOS-XE devices like Cisco Catalyst 9000 Series switches and Cisco ASR 1000 Series Routers as well to external fabrics.

## Configuration Compliance in External Fabrics

With external fabrics, any Nexus switches, Cisco IOS-XE devices, Cisco IOS XR devices, and Arista 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 Nexus Dashboard Fabric Controller, 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 Nexus Dashboard Fabric Controller, 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 Nexus Dashboard Fabric Controller and the switch has additional configuration under the same top-level command, as mentioned earlier, CC will only ensure that the intent defined in Nexus Dashboard Fabric Controller is present on the switch. When this user defined intent on Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller and deployed to the switch.
2. Additional configuration exists under **router bgp** in **Running config** that does not exist in user-defined Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller.
3. The **Pending Config** and the **Side-by-side Comparison** when the intent that was pushed earlier via Nexus Dashboard Fabric Controller is deleted from Nexus Dashboard Fabric Controller by deleting the switch_freeform policy that was created in the Step 1.
4. A **switch_freeform** policy with the top-level **router bgp** command needs to be created. This enables CC to generate the configuration needed to remove only the desired sub-config which was pushed from Nexus Dashboard Fabric Controller earlier.
5. The removed configuration is only the subset of the configuration that was pushed earlier from Nexus Dashboard Fabric Controller.

   For interfaces on the switch in the external fabric, Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller. 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:

### Procedure

**Step 1** Select **LAN** > **Fabrics**.

Double click on the fabric name to view **Fabric Overview** screen.

**Step 2** On the Switches tab, double click on the switch name to view **Switch Overview** screen.

On the Policies tab, all the policies applied on the switch within the chosen fabric are listed.

**Step 3** On the Policies tab, from the **Actions** drop-down list, select **Add Policy**.

**Step 4** Add a Policy Template Instances (PTIs) with the required configuration CLIs using the **switch_freeform** template and click **Save**.

**Step 5** Select the created policy and select **Push Config** from the **Actions** drop-down list to deploy the configuration to the switch(es).

## Managing Cisco IOS-XR Devices using NDFC

In general, workload requires communication with services outside of the data center domain in a data center fabric. This includes users accessing an application and services from the internet and WAN. VXLAN EVPN fabrics with border devices are considered as a handoff for north-south connectivity. These border devices are in peer with IOS-XR routers, which is a backbone routers for WAN and internet connectivity.

In DCNM Release 11.5(x), users with an admin role can control VXLAN EVPN fabrics with capabilities such as monitoring, automation, and compliance. You can only monitor the IOS-XR routers in monitored mode. Therefore, there is a requirement for a single fabric controller to manage, and automate configurations between these devices to balance and check configurations compliance for communicating between different services.

From NDFC Release 12.0.1a, users with an admin role can manage IOS-XR routers which is limited to automation and checking compliance. New templates and policies are introduced to automate and manage eBGP VRF Lite handoff between border switches and IOS-XR routers. NDFC allows you to check configuration compliance for IOS-XR devices similar to Cisco Nexus switches in the external fabrics.

**Note** For all non-Nexus devices, only MD5 protocol is supported for SNMPv3 authentication.

### Configuring IOS-XR as Edge Router

To extend VRF Lite from Cisco Nexus 9000 fabric with border devices for IOS-XR as edge router, refer to *VRF Lite Between Cisco Nexus 9000 Based Border and Non-Nexus Device* section.

For more information, see video at Managing and Configuring ASR 9000 using NDFC.

## Configuring Non-Nexus Devices for Discovery

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

### *Configuring IOS-XE Devices for Discovery*

**Note** In case of failure or issues configuring devices contact Cisco Technical Assistance Center (TAC).

Before you discover the Cisco IOS-XE devices in Nexus Dashboard Fabric Controller, 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** Before you run SNMP command on the switch, ensure that the IP addresses, username and SNMP related configurations are defined on the switch. Run the following SNMP command on the switch console.

```
aaa new-model
aaa session-id common
 ip domain name cisco
username admin privilege 15 secret 0 xxxxx
snmp-server group group1 v3 auth read view1 write view1
snmp-server view view1 mib-2 included
snmp-server view view1 cisco included
snmp-server user admin group1 v3 auth md5 xxxxx priv des xxxxx
line vty 0 4
privilege level 15
transport input all
line vty 5 15
privilege level 15
transport input all
line vty 16 31
transport input ssh
```

*Configuring Arista Devices for Discovery*

Enable Privilege Exec mode using the following command:

```
switch> enable
switch#

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

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

```
switch# configure terminal
switch (config)# username ndfc privilege 15 role network-admin secret cisco123
snmp-server view view_name SNMPv2 included
snmp-server view view_name SNMPv3 included
snmp-server view view_name default included
snmp-server view view_name entity included
snmp-server view view_name if included
snmp-server view view_name iso included
```

```
snmp-server view view_name lldp included
snmp-server view view_name system included
snmp-server view sys-view default included
snmp-server view sys-view ifmib included
snmp-server view sys-view system included
snmp-server community private ro
snmp-server community public ro
snmp-server group group_name v3 auth read view_name
snmp-server user username group_name v3 auth md5 password priv aes password
```



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

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

### Show Run Command

```
switch (config)# snmp-server engineID local f5717f444ca824448b00
snmp-server view view_name SNMPv2 included
snmp-server view view_name SNMPv3 included
snmp-server view view_name default included
snmp-server view view_name entity included
snmp-server view view_name if included
snmp-server view view_name iso included
snmp-server view view_name lldp included
snmp-server view view_name system included
snmp-server view sys-view default included
snmp-server view sys-view ifmib included
snmp-server view sys-view system included
snmp-server community private ro
snmp-server community public ro
snmp-server group group_name v3 auth read view_name
snmp-server user user_name
            group_name v3 localized f5717f444ca824448b00 auth md5
be2eca3fc858b62b2128a963a2b49373 priv aes be2eca3fc858b62b2128a963a2b49373
!
spanning-tree mode mstp
!
service unsupported-transceiver labs f5047577
!
aaa authorization exec default local
!
no aaa root
!
username admin role network-admin secret sha512
$6$5ZKs/7.k2UxrWDg0$FOkdVQsBTnOquW/9AYx36YUBSPNLFdeuPIse9XgyHSdEOYXtPyT/0sMUYYdkMffuIjgn/d9rx/Do71XSbygSn/
username cvpadmin role network-admin secret sha512
$6$fLGFj/PUcuJT436i$Sj5G5c4y9cYjI/BZswjjmZW0J4npGrGqIyG3ZFk/ULza47Kz.d31q13jXA7iHM677gwqQbFSH2/3oQEaHRq08.
username ndfc privilege 15 role network-admin secret sha512
$6$M48PNrCdg2EITEdG$iiB880nvFQQlrWoZwOMzdt5EfkuCIraNqtEMRS0TJUhNKCQnJN.VDLFsLAmP7kQBo.C3ct4/.n.2eRlcP6hij/
```

### Show SNMP View Command

```
configure terminal# show snmp view
view_name SNMPv2 - included
view_name SNMPv3 - included
```

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

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

*Configuring and Verifying Cisco IOS-XR Devices for Discovery*

To configure IOS-XR devices, run the following commands on the switch console:

```
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
```

Below shown example of configuring IOS-XR device on a switch.

```
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
```

To verify IOS-XR devices, run the following command:

```
RP/0/RSP0/CPU0:ios(config)#
RP/0/RSP0/CPU0:ios(config)#show run snmp-server
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.

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

| Field | Description |
|---|---|
| Seed IP | Enter the IP address of the switch.<br>You can import more than one switch by providing the IP address range. For example: 10.10.10.40-60<br>The switches must be properly cabled and connected to the Nexus Dashboard Fabric Controller server and the switch status must be manageable. |
| Device Type | • Choose **IOS XE** from the drop-down list for adding Cisco CSR 1000v, Cisco ASR 1000 Series routers, or Cisco Catalyst 9000 Series Switches.<br>• Choose **IOS XR** from the drop-down list for adding ASR 9000 Series Routers, Cisco NCS 5500 Series Routers, IOS XR Release 6.5.3 or Cisco 8000 Series Routers.<br>**Note** To add Cisco IOS XR devices in managed mode, navigate to the **General Parameters** tab in the fabric settings and uncheck the **Fabric Monitor Mode** check box.<br>• Choose **Other** from the drop-down list for adding non-Cisco devices, like Arista switches. |
| 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 Nexus Dashboard Fabric Controller 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.**

**Note**
If you see the following error message after attempting to import the switch into the fabric:

```
Error while creating the (Seed interface) intent for basic switch configurations. Please
retry using config Save/Deploy.
```

This might be because the permissions were not set properly for the switch before you tried to import it into the fabric. Set the permissions for the switch using the procedures in Configuring IOS-XE Devices for Discovery, on page 124, then try importing the switch into the fabric again.

**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 **Fabric Overview**.

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

### What to do next

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

If you added these devices under managed mode, you can add policies too.

## Managing Non-Nexus Devices to External Fabrics

From Nexus Dashboard Fabric Controller 12.0.1a, IOS-XR is supported in managed mode.



**Note** Configuration compliance is enabled for IOS-XE and IOS-XR switches, similar to the way the Nexus switches are handled in External Fabric. For more information, see Configuration Compliance in External Fabrics, on page 91.

Nexus Dashboard Fabric Controller sends commit at the end of deployment for IOS-XR devices.

Nexus Dashboard Fabric Controller provides a few templates for IOS-XR devices. Use the **ios_xr_Ext_VRF_Lite_Jython.template** for IOS-XR switch to be an edge router to establish eBGP peering with border. This will create config for vrf, eBGP peering for the vrf and the sub-interface. Similarly, **ios_xe_Ext_VRF_Lite_Jython** can be used for IOS-XE switch to be an edge router to establish eBGP peering with border.

## Creating a vPC Setup

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

### Procedure

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

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

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

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

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

**vPC Domain tab**: Enter the vPC domain details.

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

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

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

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

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

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

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

**Step 3** Click **Save**.

The **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.

After creating a vPC pair, you can view vPC details in **vPC Overview** window.

## 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 **Deploy Config**.

**Step 4** (Optional) Click the value under the **Recalculate 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 **Deploy Config**: 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.

# IPFM Fabrics

This section describes how to configure fabrics related to IP Fabric for Media (IPFM). The IPFM fabric feature is a part of LAN fabric. To enable the IPFM fabrics feature, you must have enabled the following features on the LAN Fabric in **Settings** > **Feature Management**:

- IP Fabric for Media – Starts microservices corresponding to media controller.
- PTP Monitoring – Enable if required. However, PTP monitoring is used for IPFM though it is independent of IPFM.
- Performance Monitoring – Provides for base interface monitoring.

Beginning from Nexus Dashboard Fabric Controller version 12.0.1a, the IPFM fabric templates are of the following types:

- IPFM_Classic Fabric – Use the IPFM_Classic fabric template to bring in switches from an existing IPFM fabric. This template works like an external or LAN Classic Fabric where only basic switch configuration such as management VRF/interface, and hostname can be imported. You can set the attribute of the fabric to Read/Write or Read-only. For the Read-only fabric, enable the monitor mode. This template supports IPFM_Classic and Generic_Multicast technologies.
- Easy_Fabric_IPFM Fabric – Use the Easy_Fabric_IPFM template to create a new IPFM fabric with Easy Fabric management and build an underlay network for the IPFM fabric.

**Note** IPFM Easy Fabric supports only Greenfield deployments.

We recommend that you deploy a 3-node cluster if you've more than 35 switches in your NDFC deployment. If you are using a Virtual Nexus Dashboard Cluster before you begin, ensure that the Persistent IP address and required settings are enabled for telemetry. Refer to Cisco Nexus Dashboard Fabric Controller Deployment Guide.

For a fresh installation, you can choose either IPFM Easy Fabric or IPFM Classic Fabric, based on your requirement.

### Creating IPFM Fabrics

Perform the following procedures to create IPFM fabrics:

1. Create the required IPFM Fabric using the appropriate templates and set the parameters. For more information about IPFM_Classic template, see Creating an IPFM Classic Fabric, on page 133. For more information about Easy_Fabric_IPFM template, see Creating an IPFM Easy Fabric, on page 136.
2. Add switches to the fabric and set the switch roles (only spine and leaf are supported for IPFM Fabric). For more information about adding switches, discovering existing and new switches, assigning roles, and deploying switches, see Switches, on page 279.

**Note** IPFM Easy Fabric supports only Greenfield deployments.

3. In the **Fabric Overview** window of your fabric, choose **Recalculate Config** from the **Actions** drop-down list. Then, in the **Deploy Configuration** window, click the **Deploy** button to deploy the configuration. For more information, see Fabric Overview, on page 181.

   **IPFM Easy Fabric:** The underlay config of each switch is calculated based on the fabric settings, switch role, and switch platform.

   **IPFM Classic Fabric:** If you choose to have Nexus Dashboard Fabric Controller manage the interfaces for your fabric, perform host_port_resync/**Interface Config Resync** to complete the migration process for the switch. For more information about host port resync, see Sync up Out-of-Band Switch Interface Configurations, on page 84.

   The time taken by host port resync depends on the number of switches/interfaces to be synchronized.

   If you want to edit or delete an IPFM fabric, see Editing an IPFM Fabric, on page 144 or Deleting an IPFM Fabric, on page 144 respectively.

4. Edit the existing interfaces as required. For more information, see Editing an Interface for IPFM Fabrics, on page 148. For more information about any new logical interfaces, see Creating an Interface for IPFM Fabrics, on page 145.

## Creating an IPFM Classic Fabric

This section describes the procedure to create an IPFM classic fabric from the **IPFM_Classic** template.

### Procedure

**Step 1** In the **LAN Fabrics** window, from the **Actions** drop-down list, choose **Create Fabric**.

The **Create Fabric** window appears.

**Note**
When you log in for the first time, the **Lan Fabrics** window displays no entries for IPFM fabrics. After you create a fabric, it is displayed in the **Lan Fabrics** window.

**Step 2** In the **Create Fabric** window, enter a fabric name and click **Choose Template**.

The **Select Fabric Template** window appears.

**Step 3** Either search or scroll and choose the **IPFM_Classic** fabric template. Click **Select**.

The **Create Fabric** window displays the following elements:

**Fabric Name** - Displays the fabric name you entered.

**Pick Template** - Displays the template type that you selected. If you want to change the template, click it. The **Select Fabric Template** window appears. Repeat the current step.

**General Parameters**, **Advanced,** and **Bootstrap** tabs - Display the fabric settings for creating an IPFM classic fabric.

**Step 4** The **General Parameters** tab is displayed by default. The fields in this tab are:

**Fabric Technology** – Choose one of the following technologies from the drop-down list:

- **IPFM_Classic**
- **Generic_Multicast**

**Fabric Monitor Mode** – Select this check box to only monitor the fabric, but not deploy the configuration.

From Cisco NDFC Release 12.1.2e, you can configure and monitor both NBM active and passive VRFs. In NBM passive mode, NDFC will be involved only in the monitoring of IPFM fabric and not configuration except in setting up VRF mode as NBM passive.

**Enable NBM Passive Mode** – Check this check box to enable NBM mode to IPFM passive for default VRF.

**Note**
You cannot edit the existing fabric to change the NBM mode. You must delete and re-create fabric to change the NBM mode from active to passive or vice-versa.

**Enable Performance Monitoring** – Select this check box to monitor the performance of the fabric.

Ensure that you do not clear interface counters from the Command Line Interface of the switches. Clearing interface counters can cause the Performance Monitor to display incorrect data for traffic utilization. If you must clear the counters and the switch has both `clear counters` and `clear counters snmp` commands (not all switches have the `clear counters snmp` command), ensure that you run both the main and the SNMP commands simultaneously. For example, you must run the `clear counters interface ethernet` *`slot/port`* command followed by the `clear counters interface ethernet` *`slot/port`* `snmp` command. This can lead to a one time spike.

**Step 5** Click the **Advanced** tab. The fields in this tab are:

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

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

**Enable NDFC as Trap Host** – Select this check box to enable Nexus Dashboard Fabric Controller as a trap host.

**Enable CDP for Bootstrapped Switch** – Enables CDP on management interface.

**Inband Mgmt** – For External and Class ic LAN Fabrics, this knob enables Nexus Dashboard Fabric Controller 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 (that is, reachable over switch mgmt0 interface). The only requirement is that for Inband managed switches, there should be IP reachability from Nexus Dashboard Fabric Controller to the switches through the Nexus Dashboard data interface. 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. Nexus Dashboard Fabric Controller has a pre-check that validates that the Inband managed switch IPs are reachable over the Nexus Dashboard data interface. Once the pre-check has passed, Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller. For more information, see Inband Management in External Fabrics and LAN Classic Fabrics, on page 152.

**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 Nexus Dashboard Fabric Controller are typically bound to the eth1 or out-of-band interface. In scenarios, where the Nexus Dashboard Fabric Controller eth0/eth1 interfaces reside in the same IP subnet, the POAP services are bound to both interfaces.

**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 configurations.

**Step 6** Click the **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. When this check box is selected, automatic IP assignment for POAP is enabled.

After you enable bootstrap, you can enable the DHCP server for automatic IP address assignment for POAP using the following method:

- **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 either DHCPv4 or DHCPv6 from the 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 Dashboard Fabric Controller 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 except /64 are not supported.

If you don’t select this check box, Nexus Dashboard Fabric Controller 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 the 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 64 and 126. This field is editable if you enable IPv6 for DHCP.

**Bootstrap Freeform Config** – (Optional) Enter extra commands as needed. For example, if you are using AAA or remote authentication related configurations, you need to add these configurations in this field to

save the intent. After the devices boot up, they contain the intent defined in the **Bootstrap Freeform Config** field.

Copy-paste the running-config to a **freeform config** field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running-config. For more information about *Resolving Freeform Config Errors in Switches*, see Enabling Freeform Configurations on Fabric Switches , on page 95.

**DHCPv4/DHCPv6 Multi Subnet Scope** – Specifies the field to enter one subnet scope per line. This field is editable after you select 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.

**Step 7** Click **Save**.

The IPFM classic fabric is created and displayed in the table in the **Lan Fabrics** window.

### What to do next

After creating the fabric, perform Recalculate Config and deploy the configuration to the switches. For more information, see Fabric Overview, on page 181.

Then, edit or create an interface as appropriate. For more information, see Interface Configuration for IPFM Fabrics.

## Creating an IPFM Easy Fabric

This section describes the procedure to create an IPFM Fabric from the Easy_Fabric_IPFM fabric template.

### Procedure

**Step 1** In the **LAN Fabrics** window, from the **Actions** drop-down list, choose **Create Fabric**.

The **Create Fabric** window appears.

**Note**
When you log in for the first time, the Lan Fabrics table has no entries. After you create a fabric, it is displayed in the **Lan Fabrics** window.

**Step 2** In the **Create Fabric** window, enter a fabric name and click **Choose Template**.

The **Select Fabric Template** window appears.

**Step 3** Either search or scroll and choose the **Easy_Fabric_IPFM** template. Click **Select**.

The **Create Fabric** window displays the following elements:

**Fabric Name** - Displays the fabric name you entered.

**Pick Template** - Displays the template type that you selected. If you want to change the template, click it. The **Select Fabric Template** screen appears. Repeat the current step.

**General Parameters**, **Multicast**, **Protocols**, **Advanced**, **Manageability**, and **Bootstrap** tabs - Display the fabric settings for creating an IPFM easy fabric.

**Step 4** The **General Parameters** tab is displayed by default. The fields in this tab are:

**Fabric Interface Numbering** - Supports only numbered (point-to-point, that is, **p2p**) networks.

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

**Fabric Routing Protocol** - The IGP used in the fabric, OSPF, or IS-IS.

**Fabric Routing Loopback Id**: The loopback interface ID is populated as 0 since loopback0 is usually used for fabric underlay IGP peering purposes. The valid value ranges from 0 to 1023.

**Manual Fabric IP Address Allocation** - Select this check box to disable dynamic allocation of fabric IP address.

- By default, Nexus Dashboard Fabric Controller allocates the underlay IP address resources (for loopbacks, fabric interfaces, and so on) 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 Nexus Dashboard Fabric Controller REST API Reference Guide, Release 12.0.1a* 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.
- Changing from static to dynamic allocation keeps the current IP resource usage intact. Only future IP address allocation requests are taken from dynamic pools.

**Fabric Routing Loopback IP Range** - Specifies the range of loopback IP addresses for the protocol peering.

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

**Enable Performance Monitoring** - Select this check box to monitor the performance of the fabric.

Ensure that you do not clear interface counters from the Command Line Interface of the switches. Clearing interface counters can cause the Performance Monitor to display incorrect data for traffic utilization. If you must clear the counters and the switch has both `clear counters` and `clear counters snmp` commands (not all switches have the `clear counters snmp` command), ensure that you run both the main and the SNMP commands simultaneously. For example, you must run the `clear counters interface ethernet` *`slot/port`* command followed by the `clear counters interface ethernet` *`slot/port`* `snmp` command. This can lead to a one time spike.

**Step 5** Click the **Multicast** tab. The fields in this tab are:

From Cisco NDFC Release 12.1.2e, you can configure and monitor both NBM active and passive VRFs. In NBM passive mode, NDFC will be involved only in the monitoring of IPFM fabric and not configuration except in setting up VRF mode as NBM passive.

**Note**
You cannot deploy VRF on switch in ROM.

**Enable NBM Passive Mode** - Select this check box to enable NBM mode to pim-passive. If you enable NBM passive mode, the switch ignores all RP and MSDP configurations. This is a mandatory check box. If you

select this check box, the remaining fields and check boxes are disabled. For more information, refer to the Configuring an NBM VRF for Static Flow Provisioning section of the *Cisco Nexus 9000 Series NX-OS IP Fabric for Media Solution Guide, Release 10.2(x)*.

You must add the **IP PIM Passive** command when you add VRF which is in passive mode to the interface. Perform the below steps to add the **IP PIM Passive** command:

- On the **Fabric Overview** window, choose **Links** > **Links** .
- Select the appropriate fabric with the policy **int_ipfm_intra_fabric_num_link** and choose **Actions** > **Edit**.

  The **Link Management - Edit Link** window appears.
- On the **General Parameters** tab, enter default or default VRF for the **Interface VRF** name.
- Click the **Advanced** tab, enter **IP PIM Passive** on the **Source Interface Freeform Config** and **Destination Interface Freeform Config** fields.
- Click **Save**.

You cannot edit the existing fabric to change the NBM mode. You must delete and re-create fabric to change the NBM mode from active to passive or vice-versa.

**Enable ASM** - Select this check box to enable groups with receivers sending (*,G) joins. If you select this check box, the ASM-related section is enabled.

**NBM Flow ASM Groups for default VRF (w/wo SPT-Threshold Infinity)** - This section comprises ASM-related information.

- Click the expander arrow next to the title of this section to collapse or expand the section.
- Use the **Actions** drop-down list to add, edit, or delete the ASM groups in the table.
  - **Add** - Choose this option to open the **Add Item** window. In the **Add Item** window, perform the following steps:
    a. Enter the appropriate values in the fields and check or clear the check box as follows:
       - **Group_Address** - Specify the IP address for the NBM flow ASM group subnet.
       - **Prefix** - Specify the subnet mask length for the ASM group subnet. The valid value for the subnet mask length ranges from 4 to 32. For example, 239.1.1.0/25 is the group address with the prefix.
       - **Enable_SPT_Threshold** - Check this check box to enable SPT threshold infinity.

    b. Click **Save** to add the configured NBM flow ASM groups to the table or click **Cancel** to discard the values.
  - **Edit** - Select the check box next to the group address and then choose this option to open the **Edit Item** window. Open the edit item and edit the ASM group parameters. Click **Save** to update the values in the table or click **Cancel** to discard the values.
  - **Delete** - Select the check box next to the group address and then choose this option to delete the ASM group from the table.
- The table displays the values for group address, prefix, and enable SPT threshold.

**RP Loopback Id** - The loopback ID used for the rendezvous point (RP), for multicast protocol peering purposes in the fabric underlay. The valid values range from 0 to 1023.

**Fabric RP Loopback IP Range** - Specifies the RP Loopback IP address range.

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

**Fabric Routing Protocol Tag** - Specifies the routing process tag for the fabric.

**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 **Fabric Routing Protocol** field in the **General Parameters** tab.

**Enable OSPF Authentication** - Select the check box to enable OSPF authentication. Clear 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 the Retrieving the Authentication Key , on page 142 section for details.

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

**Enable IS-IS Network Point-to-Point** - Select the check box to enable network point-to-point on fabric interfaces which are numbered.

**Enable IS-IS Authentication** - Select the check box to enable IS-IS authentication. Clear 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, for example, CiscoisisAuth.

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

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

**Note**
Plain text passwords are not supported. Log in to the switch, retrieve the encrypted key and enter it in this field. Refer the Retrieving the Authentication Key , on page 142 section for details.

**Enable PIM Hello Authentication** - Enables the PIM hello authentication.

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

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

**Intra Fabric Interface MTU** - Specifies the MTU for the intra fabric interface. This value must be an even number. The valid values range from 576 to 9216. This is a mandatory field.

**Layer 2 Host Interface MTU** - Specifies the MTU for the layer 2 host interface. This value must be an even number. The valid values range from 1500 to 9216.

**Power Supply Mode** - Choose the appropriate power supply mode that will be the default mode for the fabric from the drop-down list. This is a mandatory field.

**Enable CDP for Bootstrapped Switch** - Select this check box to enable CDP on management (mgmt0) interface for bootstrapped switch. By default, for bootstrapped switches, CDP is disabled on the mgmt0 interface.

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

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

**Enable Precision Time Protocol (PTP)** - Enables PTP across a fabric. When you select this check box, PTP is enabled globally and on intra-fabric 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 79.

**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 loopback ID. Otherwise, an error appears. The PTP loopback ID can be the same as BGP loopback or user-defined loopback which is created from Nexus Dashboard Fabric Controller. The PTP loopback will be created automatically if it is not created.

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

**PTP Profile** - Select a PTP profile from the list. PTP profile is enabled only on ISL links. The supported PTP Profiles are IEEE-1588v2, SMPTE-2059-2, and AES67-2015.

**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.

**Step 8** Click the **Manageability** tab. The fields in this tab are:

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

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

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

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

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

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

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

**AAA Freeform Config** - Specifies the AAA freeform Configurations.

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

**Step 9** Click the **Bootstrap** tab. The fields in this tab are:

**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 for POAP 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, Nexus Dashboard Fabric Controller 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** field is disabled.

**Note**
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 except /64 are not supported.

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

**DHCP Scope End Address** - Specifies the last IP address in 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 64 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 about *Resolving Freeform Config Errors in Switches*, see Enabling Freeform Configurations on Fabric Switches , on page 95.

**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

**Step 10** Click **Save**.

The Easy Fabric IPFM is created and displayed in the table in the **Lan Fabrics** window.

### What to do next

After creating the fabric, perform Recalculate Config and deploy the configuration to the switches. For more information, see Fabric Overview, on page 181.

Then, edit or create an interface as appropriate. For more information, see Interface Configuration for IPFM Fabrics.

## 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.

## Editing an IPFM Fabric

In the **LAN Fabrics** window, select the fabric that you want to edit. From the **Actions** drop-down list, choose **Edit Fabric**. Edit the fields in the template as required. Click **Save**.



**Note** After the fabric settings are changed, perform Recalculate Config, and deploy the configuration to the switches.

## Deleting an IPFM Fabric

In the **LAN Fabrics** window, select the fabric that you want to delete. From the **Actions** drop-down list, choose **Delete Fabric**. When a message appears asking whether you want to delete the fabric, click **Confirm**.

## Interface Configuration for IPFM Fabrics

Cisco Nexus Dashboard Fabric Controller Web UI allows you to configure IPFM External-Links for each switch in your fabric. The external device can connect to the network through this interface by marking it as IPFM External-Link.



**Note** A user with the network operator role in Nexus Dashboard Fabric Controller cannot save, deploy, undeploy, or edit interface configs.

Beginning with NDFC Release 12.0.1a, Interfaces in IPFM fabrics are managed by the Nexus Dashboard Fabric Controller Interface Manager. The default interface policy for IPFM is **int_ipfm_l3_port**.

The following issues are seen when NBM VRF is deleted from NDFC after interface is enabled with NBM external-link and unicast BW setting. When this occurs, the affected interfaces continues to show external-link and ucast BW as set. Perform the following steps to cleanup:

1. Select all the switches that has these interface issues under **Policies** tab using **Add Policy**.
2. Choose **host_port_resync** template and click **Save**.
3. Select **Recalculate & Deploy**. This syncs switch configuration with NDFC.
4. Select **Resync All**.

The non-fabric ethernet interface policy templates for IPFM fabrics are **int_ipfm_l3_port**, **int_ipfm_access_host**, and **int_ipfm_trunk_host**.

The port channel interface policy templates for IPFM fabrics are **int_ipfm_port_channel_access_host**, **int_ipfm_port_channel_trunk_host**, **int_ipfm_port_channel_access_member**, and **int_ipfm_port_channel_trunk_member**.

The Switch Virtual Interface (SVI) template for IPFM fabrics is **int_ipfm_vlan**.

## Creating an Interface for IPFM Fabrics

This section describes the procedure to create a new interface for an IPFM fabric based on the template that you have selected from the available IPFM fabric interface templates.



**Note** IPFM fabrics do not support V6 underlay.

### Procedure

**Step 1** Navigate to the **Fabric Overview** window for your fabric and click the **Interfaces** tab.

**Step 2** Choose **Create new interface** from the **Actions** drop-down list.

The **Create new interface** window appears.

**Step 3** Select either Port Channel, Loopback, or SVI as the interface type for IPFM.

**Step 4** Select a device from the drop-down list. The switches (spine and leaf) that are a part of the fabric are displayed in the drop-down list.

**Step 5** Enter the Port Channel ID, Loopback ID, or VLAN ID, based on your choice of the interface type.

**Step 6** Click the **No Policy Selected** link to select a policy that is specific to IPFM. In the **Select Attached Policy Template** dialog box, choose the required interface policy template and click **Save**.

**Step 7** Enter the appropriate values in the **Policy Options** area. Note that the appropriate Policy Options fields are displayed based on the policy.

- **Type - Port Channel**

  **Port Channel Member Interfaces** - Specify a list of member interfaces, for example, e1/5,eth1/7-9.

  **Port Channel Mode** - Select one of the following channel mode options: on, active, or passive.

**Enable BPDU Guard** - Select one of the following options for spanning-tree Bridge Protocol Data Unit (BPDU) guard:

- true - enables bdpuguard
- false - disables bpduguard
- no - returns to default settings

**Enable Port Type Fast** - Select this check box to enable spanning-tree edge port behavior.

**MTU** - Specify the maximum transmission unit (MTU) for the Port Channel or the MTU for the interface. The valid value range for MTU for the interface is from 576 to 9216.

**SPEED** - Specify the port channel speed or the interface speed.

**Access Vlan** - Specify the VLAN for the access port.

**Trunk Allowed Vlans** - Enter one of the following values:

- none
- all
- vlan ranges, for example, 1-200, 500-2000, 3000)

**Enable PTP** - Select this check box to enable Precision Time Protocol (PTP) for the host interface for the IPFM fabric. For more information about PTP, see PTP Configuration for IPFM Fabrics, on page 147.

**PTP Profile** - Select a PTP profile from the drop-down list: **IEEE-1588v2**, **SMPTE-2059-2**, or **AES67-2015**.

**PTP Vlan** - Specifies the PTP vlan for member interface when PTP is enabled.

**Port Channel Description** - Enter description for the port channel.

**Freeform Config** - Enter additional CLI for the port channel if required.

**Enable Port Channel** - Select this check box to enable the port channel.

- **Type - Loopback**

  **Interface VRF** - Enter the name of the interface VRF. Enter **default** for default VRF.

  **Loopback IP** - Enter an IPv4 address for the loopback interface.

  **Loopback IPv6 address** - Enter an IPv6 address for the loopback interface if the VRF is non-default. For default VRF add the IPv6 address in the freeform.

  **Route-Map TAG** - Enter the Route-Map tag associated with the interface IP.

  **Interface Description** - Enter description for the interface. The maximum size limit is 254 characters.

  **Freeform Config** - Enter additional CLI for the loopback interface if required.

  **Enable Interface** - Select this check box to enable the interface.

- **Type - SVI**

  **Interface VRF** - Enter the name of the interface VRF. Enter **default** for default VRF.

  **VLAN Interface IP** - Enter IP address of the VLAN interface.

**IP Netmask Length** - Specify the IP netmask length used with the IP address. The valid value range is from 1 to 31.

**Routing TAG** - Enter the routing tag associated with the interface IP.

**MTU** - Specify the maximum transmission unit (MTU) for the Port Channel or the MTU for the interface. The valid value range for MTU for the interface is from 576 to 9216.

**Disable IP redirects** - Select this check box to disable both IPv4 and IPv6 redirects on the interface.

**IPFM External-Link** - Select this check box to specify that the interface is connected to an external router.

**Interface Description** - Enter description for the interface. The maximum size limit is 254 characters.

**Freeform Config** - Enter additional CLI for the VLAN interface if required.

**Interface Admin State** - Select this check box to enable admin state for the interface.

**Step 8** Based on your requirements, click one of the following buttons:

- Save - Click **Save** to save the configuration changes.
- Preview - Click **Preview** to open the **Preview interfaces configuration** window and view the details.
- Deploy - Click **Deploy** to configure the interfaces.

### What to do next

If you want to edit the interface, see Editing an Interface for IPFM Fabrics, on page 148.

If your interface is ready, add a policy for configuring the IPFM fabric. For more information, see Adding a Policy for Configuring an IPFM Fabric, on page 148

## *PTP Configuration for IPFM Fabrics*

The Precision Time Protocol (PTP) is a protocol used to synchronize clocks throughout a computer network. When creating an interface, if you enable the **Enable PTP** check box, PTP is enabled across the fabric and on all the intrafabric interfaces. The supported PTP profiles for IPFM fabrics are **IEEE-1588v2**, **SMPTE-2059-2**, and **AES67-2015**.

A few things to note about the per-interface PTP profile for nonfabric ethernet interfaces are as follows:

- You must enable PTP and select PTP profile on each nonfabric ethernet interface.
- PTP profile can be different from the fabric level one.
- PTP must be enabled in the fabric settings before PTP can be configured on a nonfabric ethernet interface.

If PTP is disabled from the fabric settings, the PTP config will be removed from all the interfaces, that is, both the fabric and nonfabric interfaces.

For more information about PTP monitoring for IPFM fabrics, see PTP (Monitoring), on page 306.

### Editing an Interface for IPFM Fabrics

This section describes the procedure to edit an existing IPFM fabric interface template. You can either change a template or edit the values for any of the editable parameters in the **Policy Options** area.

#### Procedure

**Step 1** Navigate to the **Fabric Overview** window for your fabric and click the **Interfaces** tab.

**Step 2** Choose **Edit interface** from the **Actions** drop-down list.

The **Edit interface** window appears.

**Step 3** This step is optional. To change a policy, click the policy link and select a policy that is specific to IPFM.

In the **Select Attached Policy Template** dialog box, choose the required interface policy template and click **Save**.

**Step 4** Edit the required values in the **Policy Options** area. Note that the appropriate Policy Options fields are displayed based on the policy. For more information about the parameters, see Creating an Interface for IPFM Fabrics, on page 145.

Note that the following fields are specific to the int_ipfm_l3_port policy:

**IPFM Unicast Bandwidth Percentage** - Specifies the dedicated percentage of bandwidth to the unicast traffic. The remaining percentage is automatically reserved for the multicast traffic. If this field is left blank, Global Unicast Bandwidth reservation is used.

**IPFM External-Link** - Select this check box to specify that the interface is connected to an external router.

**Border Router** - Select this check box to enables the border router configuration on the interface. The interface is a boundary of a PIM domain.

**Interface Description** - Enter description for the interface. The maximum size limit is 254 characters.

**Step 5** Based on your requirements, click one of the following buttons:

- Save - Click **Save** to save the configuration changes.
- Preview - Click **Preview** to open the **Preview interfaces configuration** window and view the details.
- Deploy - Click **Deploy** to configure the interfaces.

#### What to do next

Add a policy for configuring the IPFM fabric. For more information, see Adding a Policy for Configuring an IPFM Fabric, on page 148.

## Adding a Policy for Configuring an IPFM Fabric

For configuration that is not uniform for all leafs or spines, additional templates are provided to help you complete the configuration of an IPFM fabric.

For example, if you enable NAT on a 9300 switch, you can create an **ipfm_tcam_nat_9300** policy to configure the required NAT TCAM for the switch.

Use the ipfm_telemetry policy for telemetry and ipfm_vrf policy for VRF config (routing, pim, asm).

**Procedure**

**Step 1** Navigate to the **Fabric Overview** window for your fabric and click the **Policies** tab.

**Step 2** Choose **Add Policy** from the **Actions** drop-down list.

The **Create Policy** window appears.

**Step 3** Click the right arrow in the **Select Switches** field.

The **Select Switches** dialog box appears.

**Step 4** Select one or more switches and click **Select**.

**Step 5** In the **Create Policy** window, click **Choose Template**.

**Step 6** In the **Select a Policy Template** dialog box, select the required template for IPFM fabric, for example, **ipfm_tcam_nat_9300**. Click **Select**.

**Step 7** Enter a priority for the template. The valid value ranges from 1 to 1000.

**Step 8** Enter the values in the TCAM-related fields. Make sure that you enter the TCAM size in increments of 256 and click **Save**.

## Editing a Policy for an IPFM Fabric

You can edit a policy for any switch in the IPFM fabric.

**Procedure**

**Step 1** Navigate to the **Fabric Overview** window for your fabric and click the **Policies** tab.

**Step 2** Search for the policy template.

**Step 3** Select the policy and choose **Edit Policy** from the **Actions** drop-down list.

The **Edit Policy** window appears.

**Step 4** Make the required changes and click **Save**.

# Netflow Support

Configuring Netflow at the fabric level allows you to collect, record, export, and monitor network flow and data to determine network traffic flow and volume for further analysis and troubleshooting. From Cisco NDFC Release 12.0.2, you can configure Netflow for Easy Fabrics, Easy Fabric eBGP, External Fabric, and LAN Classic templates.

After netflow is enabled for fabric, you can configure netflow on a network, or an interface (VLAN, SVI, physical interface, sub-interface, or port-channel). Before enabling netflow on the interface or network, ensure that the specified monitor name is defined in the fabric settings.

When Netflow is enabled at the Fabric level, the configuration is generated for netflow capable switches (FX/GX/EX) in the fabric except for spine/super-spine or switches with **no_netflow** policy. In a Multi-Site domain configuration, netflow is configured per Easy Fabric and not for the entire Multi-Site domain.

**Note** NDFC does not validate the **Netflow Monitor** name.

The following are the guidelines for Netflow configuration on other networks elements:

- For VRF Lite IFC, the netflow configuration is not inside the configuration profile, regardless of overlay mode.
- For networks, netflow configurations are not inside the configuration profile, regardless of overlay mode.
- You can configure netflow for Layer 2 Interface on trunk ports, access ports, dot1q tunnels, Layer2 port-channel, and VPC ports.
- You can configure netflow for the Layer 3 interface on SVI, Routed host, L3 Port-Channel, and sub-interfaces.
- Netflow configuration for VLANs uses **vlan_netflow** Record Template. In Brownfield deployment, the netflow configuration for VLANs is in switch freeform.
- You can enable Netflow under SVI (for routed traffic) or Vlan Configuration (for switched traffic).
- To configure IPv6 flow monitoring, use **switch_freeform** or **interface freeform**.
- Netflow configuration under the trunk or routed port is in **interface freeform**.
- For Host port resync, netflow configuration is captured in interface freeform.
- There is no explicit support for netflow in Intra-Fabric link or Multisite Underlay IFC. Note that you can use freeform configuration.

### Netflow Support for Brownfield deployments

For Brownfield deployments, global netflow configuration for export, record, and monitor are not captured due to the telemetry use case. After brownfield import, to avoid global level netflow command being removed, you can perform the following actions:

- Do not turn on strict CC.
- Include the netflow global configuration in **switch freeform**.
- Enable Netflow in the fabric setting matching with the switch configuration.

  Interface and VLAN level netflow configuration on the switch will be captured in **freeform**.
- SVI netflow config is captured in **switch_freeform** tied to the network.
- Netflow configuration for trunk or routed ports is in the **interface freeform**.
- Netflow configuration for VLANs is in the **switch_freeform**.
- The sub-interface configuration for VRF-Lite extensions is in **int_freeform**.

# Precision Time Protocol for External Fabrics and LAN Classic Fabrics

In the Fabric settings for the **External Fabric** 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.



**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 Nexus Insights for Cisco Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller using the **host_port_resync** policy. For more information, see Sync up Out-of-Band Switch Interface Configurations, on page 84.

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 **Deploy Config**. 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 Nexus Dashboard Fabric Controller. Inter-fabric links do not support PTP or ttag configurations if both fabrics are managed by Nexus Dashboard Fabric Controller.
- TTAG configuration is configured by default after the breakout. After the links are discovered and connected post breakout, perform **Deploy Config** to generate the correct configuration based on the type of port (host, intra-fabric link, or inter fabric link).

# Brownfield Deployment-Transitioning VXLAN Fabric Management to Nexus Dashboard Fabric Controller

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

# Inband Management in External Fabrics and LAN Classic Fabrics

### Inband Management

Cisco Nexus devices have dedicated out-of-band (OOB) management ports (mgmt0) to manage devices via telnet or SSH connections.

Now you can manage Cisco Nexus devices via Inband using front panel ports either by assigning management IP addresses on one of the ports or using loopback or SVI. By default, (mgmt0) interface is part of management VRF.

In NDFC by default, VRF is used for Inband management, you can use other defined VRFs for inband management for nexus devices. Inband Management is the ability to administer a network through LAN connection.

You can 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 the inband switch using seed IP. If there are multiple VRFs for the same seed IP, then no intent will be learned for seed interfaces. You must create intent or configuration manually.

After configuring or editing the Fabric settings, you must Deploy Config. You cannot change the Inband Mgmt settings after you import inband managed switches to the Fabric. If you uncheck the check box, 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 are 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.
```

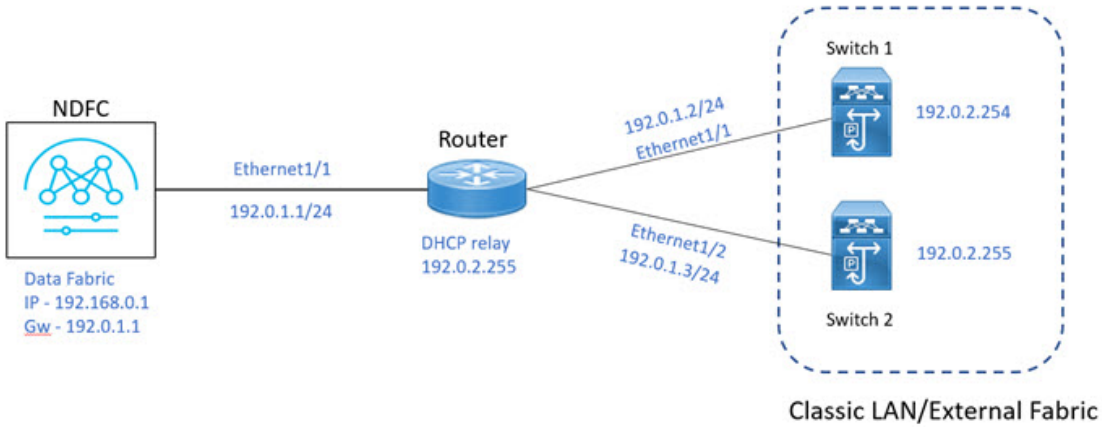
Create a policy for next-hop interfaces. Routes to Cisco Nexus Dashboard Fabric Controller from 3rd party devices can contain multiple interfaces, which are known as ECMP routes. Find the next-hop interface and create an intent for the switch. Interface IP and VRF changes are not allowed.

If inband management is enabled, during Image management, the data interface of nexus dashboard is used to copy images on the switch, in ISSU, EPLD, RPM & SMU installation flows.

If you import 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.




### Prerequisites

The following are the prerequisites for using Inband Management:

- Configure appropriate Data Network Routes for reachability to the switch Inband interfaces on Cisco Nexus Dashboard. On Nexus Dashboard, choose **Admin Console** > **Infrastructure** > **Cluster Configuration**. On **General** tab, enter route IP addresses.
- On NDFC Web UI, navigate to **Server settings > Admin** and choose **Data** from **LAN Device Management Connectivity** drop-down list to manage easy fabrics through inband management, or an error message is displayed. If you choose **Data**, ensure that the required 'Data Service IPs' are available in the Nexus Dashboard **External Service Pools** tab.

### Guidelines and Limitations

The following are the guidelines and limitations for Inband Management:

- Both Inband and out-of-band switches in the same fabric is not supported.
- When you add switches to fabric, ensure that the switches are not in maintenance mode.

# Enhanced Role-based Access Control

Starting from Cisco Nexus Dashboard Fabric Controller Release 12.0.1(a), all RBAC is in Nexus Dashboard. User-roles and access are defined from Nexus Dashboard for fabrics on NDFC.

Nexus Dashboard admin role is considered as Network-admin role in NDFC.

DCNM had five roles to perform various access and operations. If a user is access a fabric with network stage role has access to all other fabrics as a network stage role. Therefore, a username is restricted with their role in DCNM.

Cisco NDFC Release 12.0.1(a) has same five roles but you can do granular RBAC with integration of Nexus Dashboard. If a user accesses a fabric as a network stage role, the same user can access different fabric with other user role such as admin or operator role. Therefore, a user can have different access on the different fabrics in NDFC.

NDFC RBAC supports following roles:

- NDFC Access Admin
- NDFC Device Upgrade Admin
- NDFC Network Admin
- NDFC Network Operator
- NDFC Network Stager

The following table describes the user roles and their privileges in NDFC.

| Roles | Privileges |
|---|---|
| NDFC Access Admin | Read/Write<br>See |
| NDFC Device Upgrade Admin | Read/Write |
| NDFC Network Admin | Read/Write |
| NDFC Network Operator | Read |
| NDFC Network Stager | Read/Write |

The following roles are supported on DCNM for backward compatibility:

- Global-admin (mapped to network-admin)
- Server-admin (mapped to network-admin)



**Note** In any window, the actions that are restricted by the user role that is logged in are grayed out.

### NDFC Network Admin

A user with the **NDFC Network Admin** role can perform all the operations in Cisco Nexus Dashboard Fabric Controller.

You can freeze a particular fabric or all fabrics in Cisco Nexus Dashboard Fabric Controller if you are a user with the **NDFC Network Admin** role.



**Note** Make sure that the switch user role for discovery or add switches or LAN credentials for NDFC must have the network-admin role.

### NDFC Device Upgrade Admin

A user with the **NDFC Device Upgrade Admin** role can perform operations only in **Image Management** window.

See the Image Management section for more information.

### NDFC Access Admin

A user with the **NDFC Access Admin** role can perform operations only in **Interface Manager** window for all fabrics.

An NDFC access admin can perform the following actions:

- Add, edit, delete and deploy layer 2 port channels, and vPC.
- Edit host vPC, and ethernet interfaces.
- Save, preview, and deploy from management interfaces.
- Edit interfaces for LAN classic, and IPFM fabrics.

  Apart from nve, management, tunnel, subinterface, SVI, interface grouping, and loopback interfaces

However, a user with the Cisco Nexus Dashboard Fabric Controller 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 Cisco Nexus Dashboard Fabric Controller is in deployment-freeze mode.

### NDFC Network Stager

A user with the **NDFC Network Stager** role can make configuration changes on Cisco Nexus Dashboard Fabric Controller. A user with the **NDFC 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 Cisco Nexus Dashboard Fabric Controller 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.

### NDFC Network Operator

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.

The difference between a network operator and a network stager is that, 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.

### Choosing Default Authentication Domain

By default login screen on Nexus Dashboard chooses the local domain for authentication. You can change domain at login time by choosing available domains from drop-down list.

Nexus Dashboard supports local and remote authentication. The remote authentication providers for Nexus Dashboard include RADIUS, and TACACS. For more information on authentication support, refer https://www.cisco.com/c/en/us/td/docs/dcn/nd/2x/user-guide/cisco-nexus-dashboard-user-guide-211.pdf.

The following table describes RBAC comparison between DCNM and NDFC access:

| DCNM 11.5(x) | NDFC 12.0.x and 12.1.x |
|---|---|
| • User has a single role.<br>• All APIs and resources are accessed with this single role. | • User can have a different role in different Nexus Dashboard for domains.<br>• Security domain contains single Nexus Dashboard, and each Dashboard contains single NDFC Fabric. |
| A single role is associated with the user by disabling or restricting the access to options in DCNM. | A single role displays only privileged resources on the selected pa restricted access are grayed out based on security domain associa selected resource on further options on NDFC. |
| DCNM AV Pair format with shells, roles, and optional access constraints. | Nexus Dashboard AV Pair format with shells, domains. |

| DCNM 11.5(x) | NDFC 12.0.x and 12.1.x |
|---|---|
| Supported roles based on deployment type LAN, SAN, or PMN. | Supported roles such as network-admin, network-operator, device-upg-admin, network-stager, access-admin are in NDFC.<br>Support for legacy roles for backward compatibility. Nexus Dashboa admin role as network-admin of DCNM. |

The following table describes DCNM 11.5(x) AV Pair format:

| Cisco DCNM Role | RADIUS Cisco-AV-Pair Value | TACACS+ Shell Cisco-AV-Pair Value |
|---|---|---|
| Network-Operator | shell:roles = "network-operator" dcnm-access="group1 group2 group5" | cisco-av-pair=shell:roles="network-operator dcnm-access="group1 group2 group5" |
| Network-Admin | shell:roles = "network-admin" dcnm-access="group1group2 group5" | cisco-av-pair=shell:roles="network-admin" dcnm-access="group1 group2 group5" |

The following table describes NDFC 12.x AV Pair format:

| User Role | AVPair Value |
|---|---|
| NDFC Access Admin | Access-admin |
| NDFC Device Upgrade Admin | Device-upg-admin |
| NDFC Network Admin | Network-admin |
| NDFC Network Operator | Network-operator |
| NDFC Network Stager | Network-stager |

The AV pair string format differs when configuring a read/write role, read-only role, or a combination of read/write and read-only roles for a specific user. A typical string includes the domain, followed by the read-write roles separated from the read-only roles using the slash (/) character; individual roles are separated by the pipe (|) character:

```
shell:domains=<domain>/<writeRole1>|<writeRole2>/<readRole1>|<readRole2>
```

# Enhanced RBAC Use-Cases

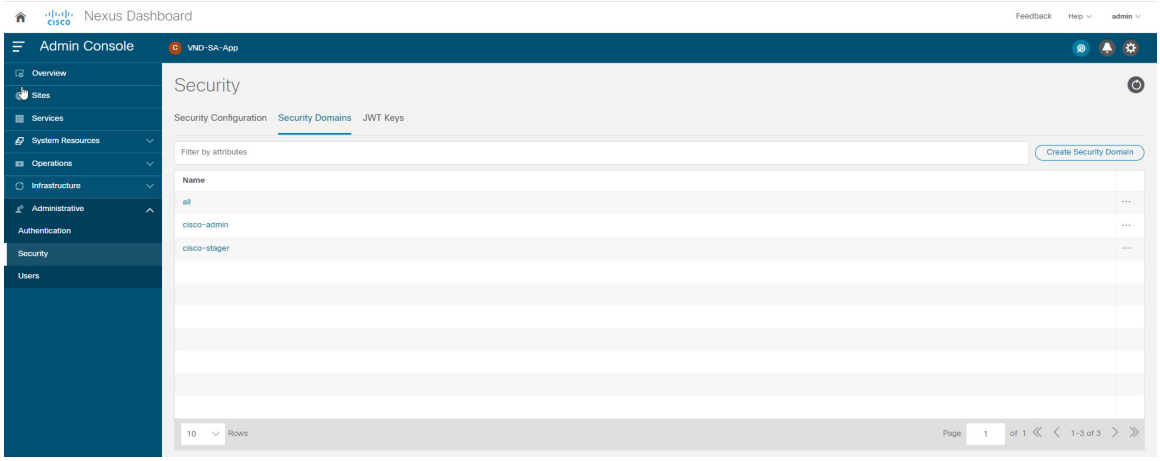
There are various fabrics in NDFC. By default a user is an admin for all the fabrics. For an example, a username **Cisco** can have admin role access to a Fabric-A and stager role access to another Fabric-B.

On Nexus Dashboard, all security policies are part of security domains. You can create the user and give access to these security domains.

To create a user and define specific roles, perform the following steps:

### Procedure

**Step 1** To create user in security domains:




a) Log in to Nexus Dashboard with admin role and navigate to **Administrative** tab.
b) On **Security Domain** tab, click **Create Security Domain** and create the following security domains:

- **all** - Similar to network-admin role. This domain has administrative access to Nexus Dashboard and NDFC service application.
- **cisco-admin** - full network-admin access to Fabric-A
- **cisco-stager** - network-stager only access to Fabric-B

**Step 2** To create a local user **Cisco**.

a) Navigate to **Users** > **Local**.
b) On **Local** tab, click **Create Local User**.

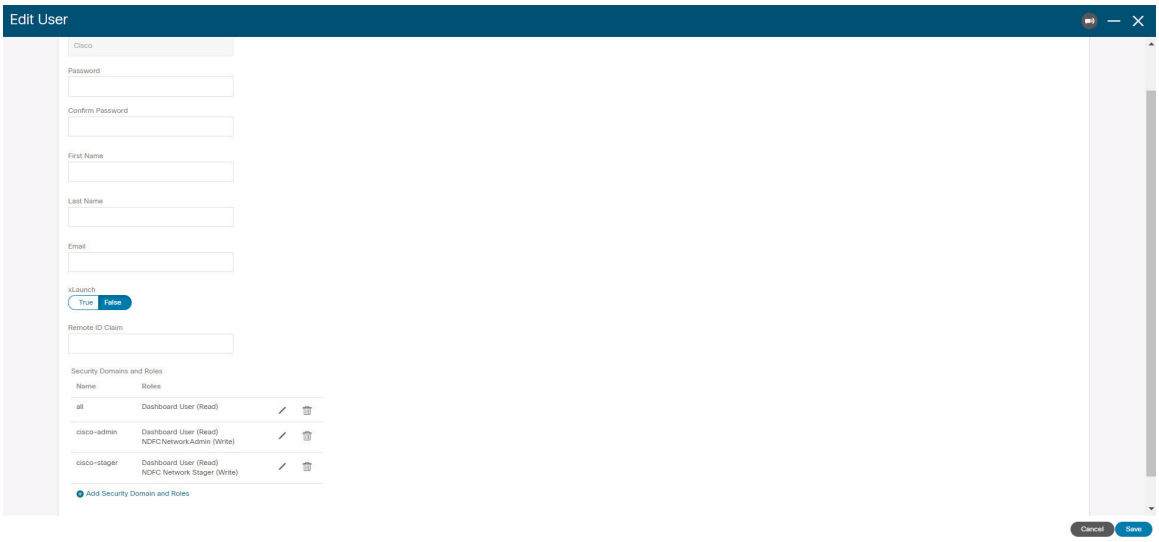
The **Create Local User** window appears.

c) Enter **Cisco** in User ID text field, provide appropriate passwords in respective fields.
d) After you create a Cisco user, navigate to **Local** window, click on **elipses** icon in **Cisco** username row and then click **Edit User**.

The **Edit User** window appears.

**Step 3** On **Edit User** window, by default, **all** security domain exists. Click **Add Security Domain** and **Roles** to add other security domains.

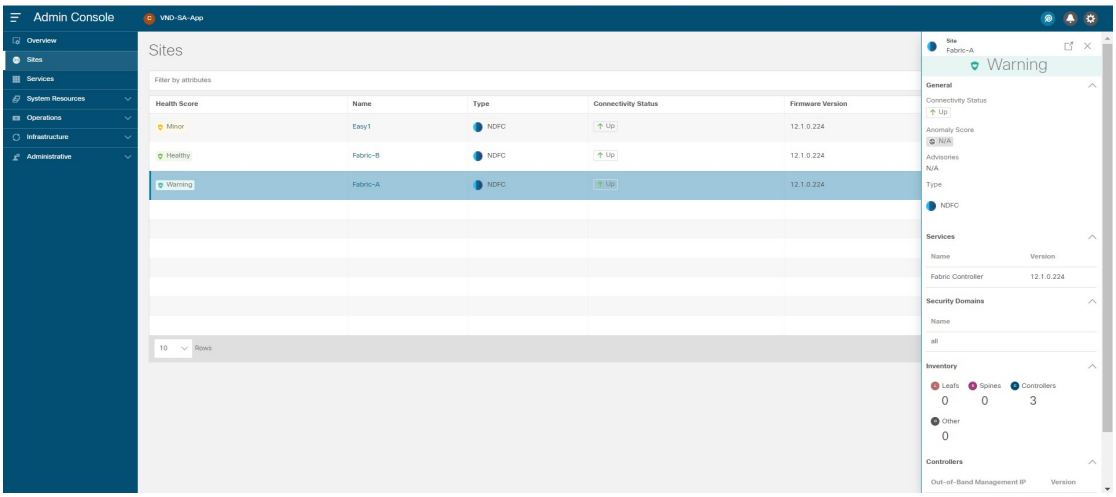
The **Add Security Domain and Roles** window appears.




**a.** Choose **cisco-admin** domain from option drop-down list and choose **NDFC Access Admin** check box and then click **Save**.

**b.** Repeat step **a** to add **cisco-stager** domain for **NDFC Network Stager** role.

**c.** To associate security domains to respective fabric sites, do the following:




On Nexus Dashboard, navigate to **Sites** window. Click on **Fabirc-A** site name.

A slide-in pane appears. You can view **all** security domain for the Fabric-A site.

**d.** To add the Cisco user as network-admin for Fabric-A, click **Elipse** icon and **Edit Site**.

**e.** Delete **all** security domain and add **network-admin** domain and save the changes.

Similarly you can add for network-stager domain.

**f.** Log out from Nexus Dashboard and log in back as **Cisco** user.

**Note**
The user role Cisco can view only NDFC related options on Nexus Dashboard based on the permissions. The user access restricted to Nexus Dashboard services.

**g.** Naviage to NDFC application.

The user Cisco can perform operations on two sites on NDFC, as the user is assigned as network-admin role for Fabric-A and network-stager role for Fabric-B.

**Note**
Network-admin role can create an interface for Fabric-A and deploy it. Whereas network-stager role can create interface for Fabric-B, but access restricted to deploy.

# Nexus Dashboard Security Domains

Access control information about a user login contains authentication data like user ID, password, and so on. Based on the authorization data, you can access resources accordingly. Admins in Cisco Nexus Dashboard can create security domains and group various resource types, resource instance, and map them into a security domain. The admins define an AV-pair for each user, which defines the access privileges for users to different resources in Cisco Nexus Dashboard. When you create a fabric, a site is created in Nexus Dashboard with the same fabric name. You can create and view these sites from **Nexus Dashboard > Sites**.

The Cisco Nexus Dashboard Fabric Controller REST APIs use this information to perform any action by checking the authorization.

**Note** When accessing REST APIs, you can verify passed payload in JSON format. Ensure that the payload is an appropriate JSON format.

When you upgrade from Cisco Nexus Dashboard Fabric Controller Release 11.x, each fabric is mapped to an autogenerated site of the same name. All these sites are mapped into the **all** security domain in Nexus Dashboard.

All resources are placed in **all** domain before they are assigned or mapped to other domains. The all security domain does not include all the available security domains in Nexus Dashboard.

### AV-Pairs

A group of security domains along with read and write roles for each domain are specified using AV-pairs. Administrators define AV-pair for each user. The AV-pair defines the access privileges to users across various resources in Nexus Dashboard.

The AV-pair format is as follows:

```
"avpair":
"shell:domains=security-domain/write-role-1|write-role-2,security-domain/write-role-1|write-role2/read-role-1|read-role-2"
```

```
For example: "avpair":
```

```
"shell:domains=all/network-admin/app-user|network-operator". "all/admin/" makes user
super-user and it's best to avoid examples with all/admin/"
```

The write role is inclusive of read role as well. Hence, `all/network-admin/` and `all/network-admin/network-admin` are the same.



**Note** From Cisco Nexus Dashboard Fabric Controller Release 12.0.1a supports the existing AV-pair format that you created in Cisco Nexus Dashboard Fabric Controller Release 11.x. However, if you are creating a new AV-pair, use the format that is mentioned above. Ensure that the shell: domains must not have any spaces.

### Specifying Cisco NX-OS User Roles and SNMPv3 Parameters on AAA Servers

You can use the VSA cisco-AV-pair on AAA servers to specify user role mapping for the Cisco NX-OS device using this format:

```
shell:roles="roleA roleB …"
```

If you do not specify the role option in the cisco-AV-pair attribute, the default user role is network-operator.

You can also specify your SNMPv3 authentication and Privacy protocol attributes as follows:

```
shell:roles="roleA roleB..." snmpv3:auth=SHA priv=AES-128
```

The SNMPv3 authentication protocol options are SHA and MD5. The Privacy protocol options are AES-128 and DES. If you do not specify these options in the cisco-AV-pair attribute, MD5 and DES are the default authentication protocols.

### Creating a Security Domain

To create a security domain from Cisco Nexus Dashboard, perform the following steps:

1. Log into Cisco Nexus Dashboard.
2. Choose **Administrative > Security**.
3. Navigate to **Security Domains** tab.
4. Click **Create Security Domain**.
5. Enter the required details and click **Create**.

### Creating a User

To create a user from Cisco Nexus Dashboard, perform the following steps:

1. Log into Cisco Nexus Dashboard.
2. Choose **Administrative > Users**.
3. Click **Create Local User**.
4. Enter the required details and click **Add Security Domain**.
5. Choose a domain from the drop-down list.
6. Assign a Cisco Nexus Dashboard Fabric Controller service read or write role by checking the appropriate check box.
7. Click **Save**.

# Backup Fabric

You can configure backup for selected fabric, from Fabric window, similarly you can configure backup on **Fabric Overview** window. Choose **Fabric Overview** > **Actions** on main window, click **Backup Fabric**.

You can back up all fabric configurations and intents automatically or manually. You can save configurations in Cisco Nexus Dashboard Fabric Controller, which are the intents. The intent may or may not be pushed on to the switches.

Cisco Nexus Dashboard Fabric Controller doesn't back up the following fabrics:

- External fabrics in monitor-only mode: 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 fabric: 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, Cisco Nexus Dashboard Fabric Controller stores all the backed-up information of the member fabrics and the MSD fabric together in a single directory.

The backed-up configuration 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.

You can enable scheduled backup for fabric configurations and intents.

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. Cisco Nexus Dashboard Fabric Controller backs up only when there's a configuration push. Cisco Nexus Dashboard Fabric Controller triggers the automatic backup only if you didn't trigger any manual backup after the last configuration push.

# Restoring Fabric

The following table describes the columns that appears on **Restore Backup** tab.

| Fields | Descriptions |
|---|---|
| Backup Date | Specifies the backup date. |
| Backup Version | Specifies the version of backup. |
| Backup Tag | Specifies the backup name. |
| NDFC Version | Specifies the version of NDFC. |
| Backup Type | Specifies the backup type, whether it is a golden backup. |

The following table describes the fields and descriptions that appears on **Action** tab.

| Actions | Descriptions |
|---|---|
| Mark as golden | To mark existing backup as golden backup, choose **Mark as golden**, a confirmation window appears, click **Confirm**. |

| Actions | Descriptions |
|---|---|
| Remove as golden | To remove existing backup from golden backup, choose **Remove as golden**, a confirmation window appears, click **Confirm**. |

To restore Fabric, perform the following procedure:

1. On Fabric Overview, select **Actions > More > Restore Fabric**.

   The **Restore Fabric** screen appears.

2. In the **Select Backup** tab, select the radio button for the backup that you choose to restore.

   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, NDFC archives only up to 10 golden backups. You can mark a backup as golden backup while restoring the fabric.

3. From the **Actions** drop-down list, select **Mark as golden** to mark the backup as Golden.

   Click **Next**.

   You can preview the details about the configuration in the backup file. You can also view the name and serial numbers for the switches in the Fabric backup.

   Click on **Delta Config** to view the configuration difference on the switches in the fabric.

4. Click **Restore Intent**.
5. On the **Restore Status** tab, you can view the status of restoring the intent.
6. Click **Next** to view the preview configuration.
7. In the **Configuration Preview** tab, you can resync the configurations on the desired switches.
8. For the desired switch, check the **Switch Name** check box, and click **ReSync**.
9. Click deploy to complete the **Restore Fabric** operation.

# VXLAN OAM

In Nexus Dashboard Fabric Controller, VXLAN OAM is supported on VXLAN Fabric, eBGP VXLAN Fabric, External, and Lan Classic fabric technologies. You can track details such as reachability and actual path of the flows in a VXLAN EVPN based-fabric topology.

**Guidelines**

- OAM must be enabled on the switches before using the OAM trace.



**Note** VXLAN OAM IPv6 is supported from Irvine release onwards.

- NX-API and NX-API on HTTP port must be enabled.
- vPC advertise-pip must be enabled.

- For switch-to-switch OAM, ensure that the VRFs are configured along with loopback interfaces with IPv4 and/or IPv6 addresses under those VRF's.
- For host-to-host OAM, ensure that the Networks are configured along with IPv4 and/or IPv6 gateway configuration.

**UI Navigation**

- In the **Topology** window: Click **Actions**. Choose **VXLAN OAM** option from the drop-down list.
- From **LAN Fabrics** window: Choose **LAN > Fabrics**. Navigate to the fabric overview window of a fabric. Click **Actions**. Choose **VXLAN OAM** option from the drop-down list.

The VXLAN OAM window appears. The **Path Trace Settings** pane on the left displays the **Switch to Switch** and **Host to Host tabs**. Nexus Dashboard Fabric Controller highlights the route on the topology between the source and destination switch for these two options.

The **Switch to Switch** option provides the VXLAN OAM ping and traceroute test results for the VTEP-to-VTEP use-case. Provide the following values to enable search by using the **Switch to Switch** option:

- In the **Source Switch** drop-down list, choose the source switch.
- In the **Destination Switch** drop-down list, choose the destination switch.
- From the **VRF** drop-down list, choose or enter the VRF details.
- Check the **All paths included** check box to include all the paths in the search results.

The **Host to Host** option provides the VXLAN OAM path trace results for the exact path that is taken by a given flow from the VTEP or switch that is connected to the source host to VTEP or switch that is connected to the destination host. For the **Host to Host** use-case, there are two options:

- VRF or SVI for a network is instantiated on the switches in the VXLAN EVPN fabric. In such a scenario, the IP address information of the end hosts is required.
- Layer 2 configuration for a given network is instantiated on the switches in the VXLAN EVPN fabric. In such a scenario, both the MAC and IP address information of the end hosts are required.

Provide the following values to enable search using the **Host to Host** option:

- From the **Source Host IP** field, enter the IPv4/IPv6 address of the source host.
- From the **Destination Host IP** field, enter the IPv4/IPv6 address of the destination host.
- In the **VRF** field, choose VRF from the drop-down list or enter the VRF name that is associated with the hosts.
- In the **Source Port** field, choose Layer 4 source port number from the drop-down list or enter its value.
- In the **Destination Port** field, choose destination port number or enter its value.
- In the **Protocol** field, choose the protocol value from the drop-down list or enter its value. This is the Layer 4 protocol, usually TCP or UDP.
- Check the **Layer 2 only** check box to search the VXLAN-EVPN fabric that is deployed in Layer 2 only mode for some networks, that is, Layer 2 VNIs. No SVIs or VRFs should be instantiated in the fabric for these networks when you use this search option. When you check this option, you have to enter details of the source MAC address, destination MAC address, and VNI too.

Click **Run Path Trace** to view the path trace from switch to switch or host to host.

You can view the forward path and reverse path as well in the topology. The summary of the path trace appears in the **Summary** tab. You can view the details of the forward and reverse paths as well under **Forward Path** or **Reverse Path** tabs. Filter the results by attributes, if needed.

# Endpoint Locator

The Endpoint Locator (EPL) feature allows real-time tracking of endpoints within a data center. The tracking includes tracing the network life history of an endpoint and getting insights into the trends that are associated with endpoint additions, removals, moves, and so on. An endpoint is anything with at least one IP address (IPv4 and\or IPv6) and MAC address. EPL feature is also capable of displaying MAC-Only endpoints. By default, MAC-Only endpoints are not displayed. An endpoint can be a virtual machine (VM), container, bare-metal server, service appliance and so on.



**Note**

- EPL is supported for VXLAN BGP EVPN fabric deployments only in the Nexus Dashboard Fabric Controller LAN fabric installation mode. The VXLAN BGP EVPN fabric can be deployed as Easy fabric, Easy eBGP fabric, or an External fabric (managed or monitored mode). EPL is not supported for 3-tier access-aggregation-core based network deployments.
- EPL displays endpoints that have at least one IP address (IPv4 and/or IPv6). EPL is also capable of displaying MAC-Only endpoints. Select the **Process MAC-Only Advertisements** checkbox while configuring EPL to enable processing of EVPN Route-type 2 advertisements having a MAC address only. L2VNI:MAC is the unique endpoint identifier for all such endpoints. EPL can now track endpoints in Layer-2 only network deployments where the Layer-3 gateway is on a firewall, load-balancer, or other such nodes.

EPL relies on BGP updates to track endpoint information. Hence, typically the Nexus Dashboard Fabric Controller must peer with the BGP Route-Reflector (RR) to get these updates. For this purpose, IP reachability from the Nexus Dashboard Fabric Controller to the RR is required. This can be achieved over in-band network connection to the Nexus Dashboard Fabric Controller Data Network interface. There is no option to configure static routes for pods on ND, so the selected RRs must be reachable through the default data network gateway.

Some key highlights of the Endpoint Locator are:

- Support for dual-homed and dual-stacked (IPv4 + IPv6) endpoints
- Support for up to two BGP Route Reflectors or Route Servers
- Support real-time and historical search for all endpoints across various search filters such as VRF, Network, Layer-2 VNI, Layer-3 VNI, Switch, IP, MAC, port, VLAN, and so on.
- Support for real-time and historical dashboards for insights such as endpoint lifetime, network, endpoint, VRF daily views, and operational heat map.
- Support for iBGP and eBGP based VXLAN EVPN fabrics. The fabrics may be created as Easy Fabrics or External Fabrics. EPL can be enabled with an option to automatically configure the spine or RRs with the appropriate BGP configuration.
- You can enable the EPL feature for upto 4 fabrics.
- EPL is supported on Multi-Site Domain (MSD).

- IPv6 underlay is not supported.
- Support for high availability
- Support for endpoint data that is stored for up to 60 days, amounting to a maximum of 2 GB storage space.
- Support for optional flush of the endpoint data to start afresh.
- Supported scale: Maximum of 50K unique endpoints per fabric. A maximum of 4 fabrics is supported. However, the maximum total number of endpoints across all fabrics should not exceed 100K.

  If the total number of endpoints across all fabrics exceeds 100K, an alarm is generated and is listed under the **Alarms** icon at the top right of the window. This icon starts flashing whenever a new alarm is generated.
- From NDFC Release 12.0.1a, Persistent or External IP addresses are required to enable EPL. For each VXLAN fabric, a specific container is spawned running a BGP instance to peer with the spines of the fabric. This container must have a persistent IP associated that is then configured as a iBGP neighbor on the spines. A different container is used for each fabric, so the number of fabrics that are managed by NDFC where EPL is enabled decides how many persistent IP addresses must be distributed for EPL. Also, the EPL establishes iBGP sessions only over the Cisco Nexus Dashboard Data interface.
- For Virtual Cisco Nexus Dashboard deployments, enable or accept promiscuous mode on the port-groups that are associated with the Nexus Dashboard Management and/or Data vNICs where IP stickiness is required. The Persistent IP addresses are given to the pods (for example, SNMP Trap/Syslog receiver, Endpoint Locator instance per Fabric, SAN Insights receiver, and so on). Every POD in Kubernetes can have multiple virtual interfaces. Specifically for IP stickiness an extra virtual interface is associated with the POD that is allocated an appropriate free IP from the external service IP pool. The vNIC has its own unique MAC address that is different from the MAC addresses associated with the vND virtual vNICs. Moreover, all communication to and from the PODs toward an external switch goes out of the same bond interface for North-to-South traffic flows. The EPL container uses Nexus Dashboard Data Interface. The Data vNIC maps to bond0 (also known as bond0br) interface. By default, the VMware system checks if the traffic flows out of a particular vNIC is matched with the Source-MAC that is associated with the vNIC. In case of NDFC, the traffic flows are sourced with the Persistent IP addresses of the given PODs. Therefore, we must enable the required settings on the VMware side.

If you are using a Virtual Cisco Nexus Dashboard Cluster before you begin, ensure that the Persistent IP addresses, EPL feature, and required settings are enabled. See below links:

Cisco Nexus Dashboard Fabric Controller Deployment Guide

Cisco Nexus Dashboard Fabric Controller Installation and Upgrade Guide

### Backup and Restore

EPL only backups data for fabrics that EPL has been configured. If EPL is disabled for a fabric(even if EPL has previously been configured there), then you cannot backup the data for that fabric. Also, you can backup only historical data (data on the Endpoint Search page).

If a backup is initiated when EPL is enabled, then when restoring the backup, the same external data IPs that EPL was using must be available on ND. If those IPs are not available, then select the **Ignore External Service IP Configuration** option in the restore backup form. However, there are chances that the EPL pods will be brought up with different IPs, so any existing EPL policies become invalid. If EPL was previously configured with the **Configure My Fabric** option, you need to disable and enable EPL so that the old policy is cleaned

up and an updated policy is deployed. If you did not use the **Configure My Fabric** option, then manually update their config with the new IPs.

# EPL Connectivity Options

Sample topologies for the various EPL connectivity options are as given below.

### NDFC Cluster Mode: Physical Server to VM Mapping

Refer to Cisco Nexus Dashboard Fabric Controller Verified Scalability Guide for more information.

# Configuring Endpoint Locator

The Nexus Dashboard Fabric Controller OVA or the ISO installation comes with two interfaces:

- Management
- Data

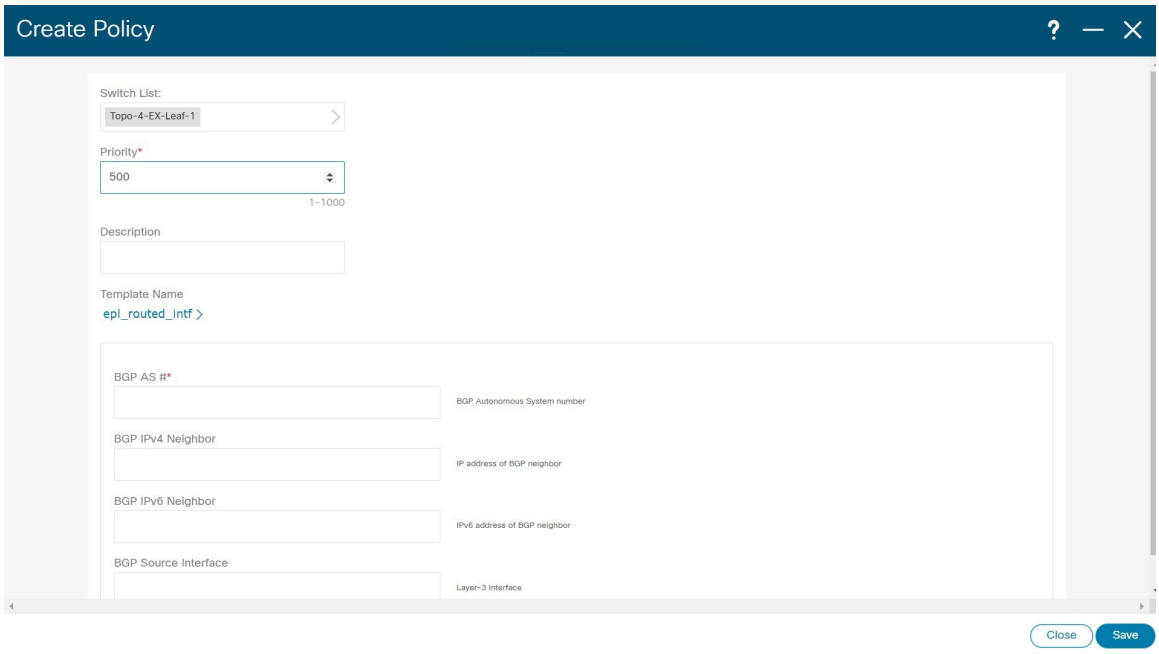
(Out-of-band or OOO) connectivity of switches via switch mgmt0 interface can be through data or Management interface. For more information refer to NDFC Installation and Upgrade Guide.

The Management interface provides reachability to the devices via the mgmt0 interface either Layer-2 or Layer-3 adjacent. This allows Nexus Dashboard Fabric Controller to manage and monitor these devices including POAP. EPL requires BGP peering between the Nexus Dashboard Fabric Controller and the Route-Reflector. Since the BGP process on Nexus devices typically runs on the default VRF, in-band IP connectivity from the Nexus Dashboard Fabric Controller to the fabric is required. The data network interface can be configured during Nexus Dashboard installation. You can't modify the configured in-band network configurations.

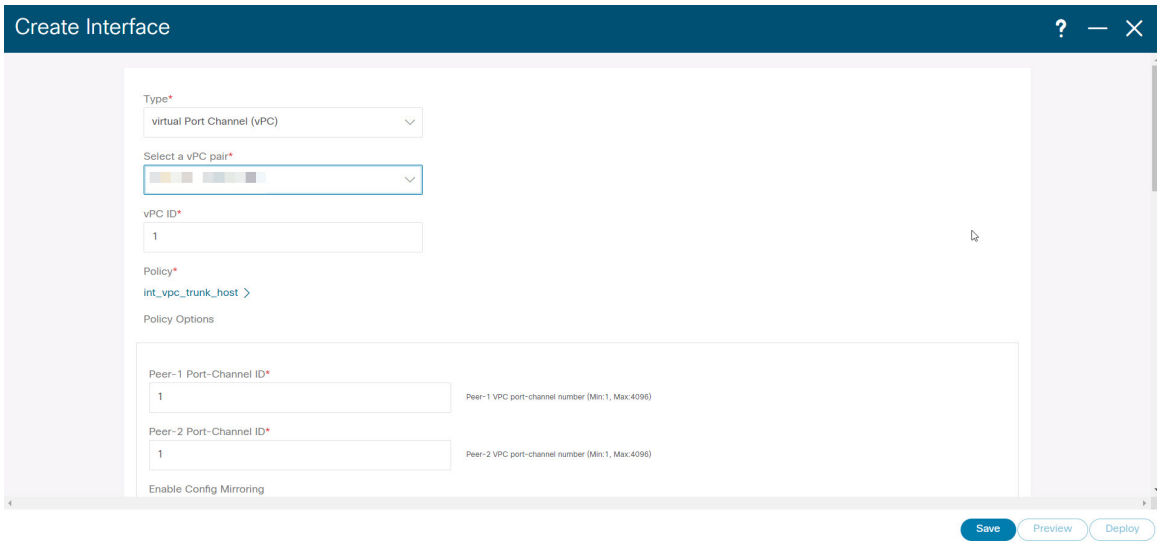


**Note** The setup of Data network interface on the Nexus Dashboard Fabric Controller is a prerequisite of any application that requires the in-band connectivity to the devices within fabric. This includes EPL and Network Insights Resources (NIR).

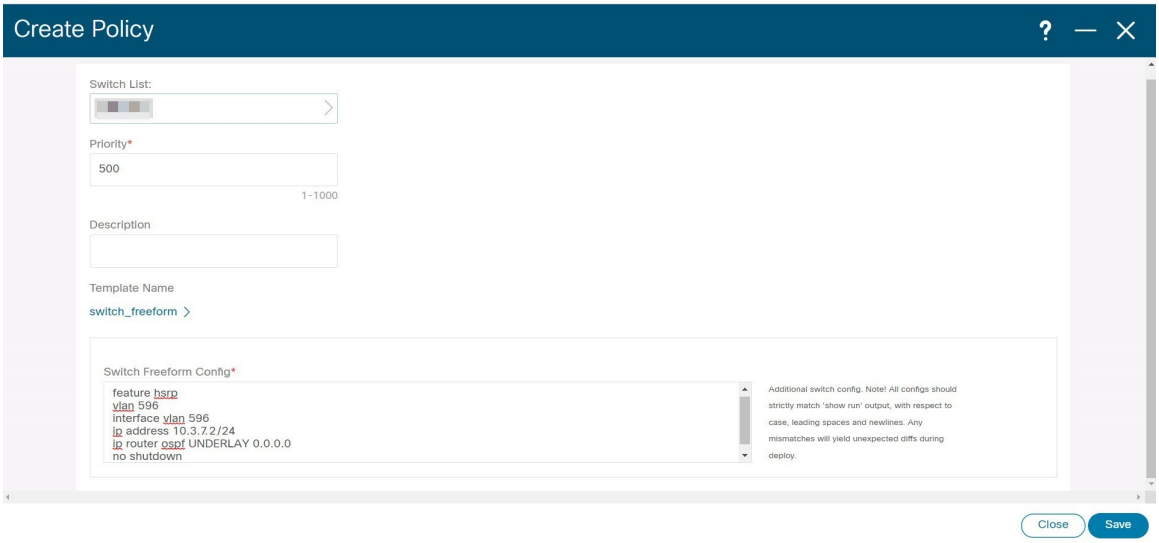
On the fabric side, for a standalone Nexus Dashboard Fabric Controller deployment, if the Nexus Dashboard data network port is directly connected to one of the front-end interfaces on a leaf, then that interface can be configured using the **epl_routed_intf** template. An example scenario of how this can be done when IS-IS or OSPF is employed as the IGP in the fabric, is depicted below:




However, for redundancy purposes, it is always advisable to have the server on which the Nexus Dashboard Fabric Controller is installed to be dual-homed or dual-attached. With the OVA Nexus Dashboard Fabric Controller deployment, the server can be connected to the switches via a port-channel. This provides link-level redundancy. To also have node-level redundancy on the network side, the server may be attached to a vPC pair of Leaf switches. In this scenario, the switches must be configured such that the HSRP VIP serves as the default gateway of the Data Network interface on the Nexus Dashboard Fabric Controller.



For the HSRP configuration on terry-leaf3, the **switch_freeform** policy may be employed as shown in the following image:




You can deploy a similar configuration on terry-leaf3 while using IP address 10.3.7.2/24 for SVI 596. This establishes an in-band connectivity from the Nexus Dashboard Fabric Controller to the fabrics over the Data Network interface with the default gateway set to 10.3.7.3.

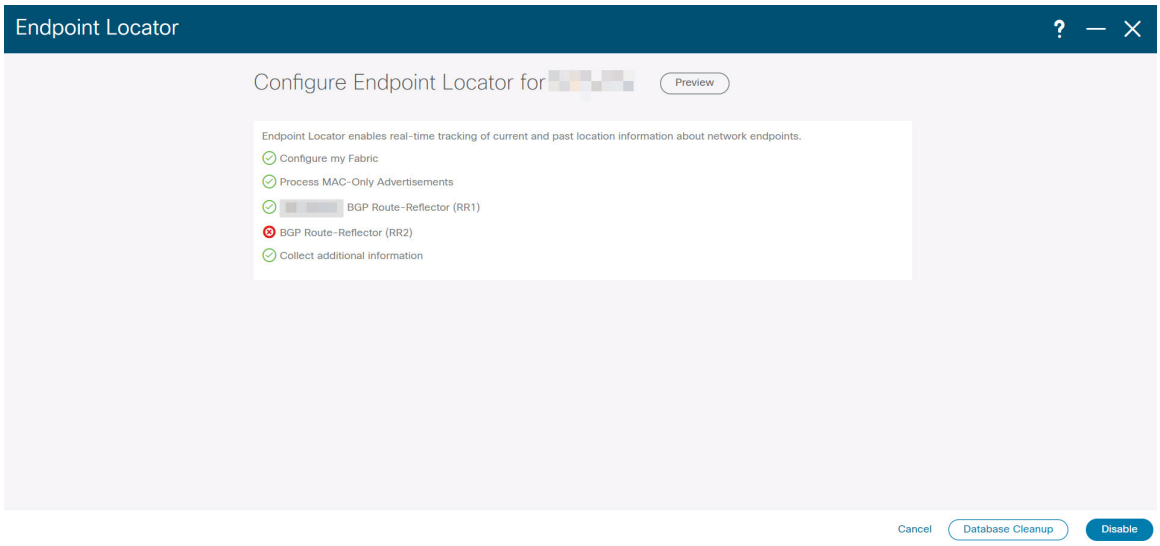
After you establish the in-band connectivity between the physical or virtual Nexus Dashboard Fabric Controller and the fabric, you can establish BGP peering.

During the EPL configuration, the route reflectors (RRs) are configured to accept Nexus Dashboard Fabric Controller as a BGP peer. During the same configuration, the Nexus Dashboard Fabric Controller is also configured by adding routes to the BGP loopback IP on the spines/RRs via the Data Network Interface gateway.

**Note** Ensure that you have enabled EPL feature for Cisco Nexus Dashboard Fabric Controller. Choose **Settings** > **Feature Management** > **Fabric Controller** choose check box **Endpoint Locator**. You can view the added EPL details on dashboard.

**Note** Cisco Nexus Dashboard Fabric Controller queries the BGP RR to glean information for establishment of the peering, like ASN, RR, IP, and so on.

To configure Endpoint Locator from the Cisco Nexus Dashboard Fabric Controller Web UI, On Fabric Overview page, choose **Actions** > **More** > **Configure Endpoint Locator**. Similarly, you can configure EPL on Topology page, right-click on required fabric, click **More** > **Configure Endpoint Locator**. The **Endpoint Locator** window appears.




You can enable EPL for one fabric at a time.

Select the switches on the fabric hosting the RRs from the drop-down list. Cisco Nexus Dashboard Fabric Controller will peer with the RRs.

By default, the **Configure My Fabric** option is selected. This option only configures EPL as a BGP neighbor of the switch and this option does not configure network reachability between EPL and the switch. This knob controls whether BGP configuration will be pushed to the selected spines/RRs as part of the enablement of the EPL feature. If the spine/RR needs to be configured manually with a custom policy for the EPL BGP neighborship, then this option should be unchecked. For external fabrics that are only monitored and not configured by Nexus Dashboard Fabric Controller, this option is greyed out as these fabrics are not configured by Nexus Dashboard Fabric Controller.

Select the **Process MAC-Only Advertisements** option to enable processing of MAC-Only advertisements while configuring the EPL feature.

**Note** If EPL is enabled on a fabric with or without selecting the **Process Mac-Only Advertisements** checkbox and you want to toggle this selection later, then you have to first disable EPL and then click **Database Clean-up** to delete endpoint data before re-enabling EPL with the desired **Process Mac-Only Advertisements** setting.

Select **Yes** under **Collect Additional Information** to enable collection of additional information such as PORT, VLAN, VRF etc. while enabling the EPL feature. To gather additional information, NX-API must be supported and enabled on the switches, ToRs, and leafs. If the **No** option is selected, this information will not be collected and reported by EPL.

**Note** For all fabrics except external fabrics, NX-API is enabled by default. For external fabrics, you have to enable NX-API in the external fabric settings by selecting the **Enable NX-API** checkbox in the **Advanced** tab of the External_Fabric_11_1 fabric template.

Click the **i** icon to view a template of the configuration that is pushed to the switches while enabling EPL. This configuration can be copied and pasted on spines or border gateway devices to enable EPL on external monitored fabrics.

Once the appropriate selections are made and various inputs have been reviewed, click **Submit** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled.

The Nexus Dashboard Data Service IP is used as BGP neighbor.

When the Endpoint Locator feature is enabled, there are a number of steps that occur in the background. Nexus Dashboard Fabric Controller contacts the selected RRs and determines the ASN. It also determines the interface IP that is bound to the BGP process. Also, appropriate BGP neighbor statements are added on the RRs or spines in case of eBGP underlay, to get them ready to accept the BGP connection that will be initiated from the Nexus Dashboard Fabric Controller. The external Nexus Dashboard Data Service IP address that is assigned to the EPL pod will be added as the BGP neighbor. Once EPL is successfully enabled, the user is automatically redirected to the EPL dashboard that depicts operational and exploratory insights into the endpoints that are present in the fabric.

For more information about the EPL dashboard, refer Monitoring Endpoint Locator, on page 180.

### Flushing the Endpoint Database

After you enable the Endpoint Locator feature, you can clean up or flush all the Endpoint information. This allows starting from a clean-slate with respect to ensuring no stale information about any endpoint is present in the database. After the database is clean, the BGP client re-populates all the endpoint information learnt from the BGP RR. You can flush the endpoint database even if you have not re-enabled the EPL feature on a fabric on which the EPL feature was previously disabled.

To flush all the Endpoint Locator information from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

#### Procedure

**Step 1** Choose **Endpoint Locator > Configure**, and click **Database Clean-Up**.

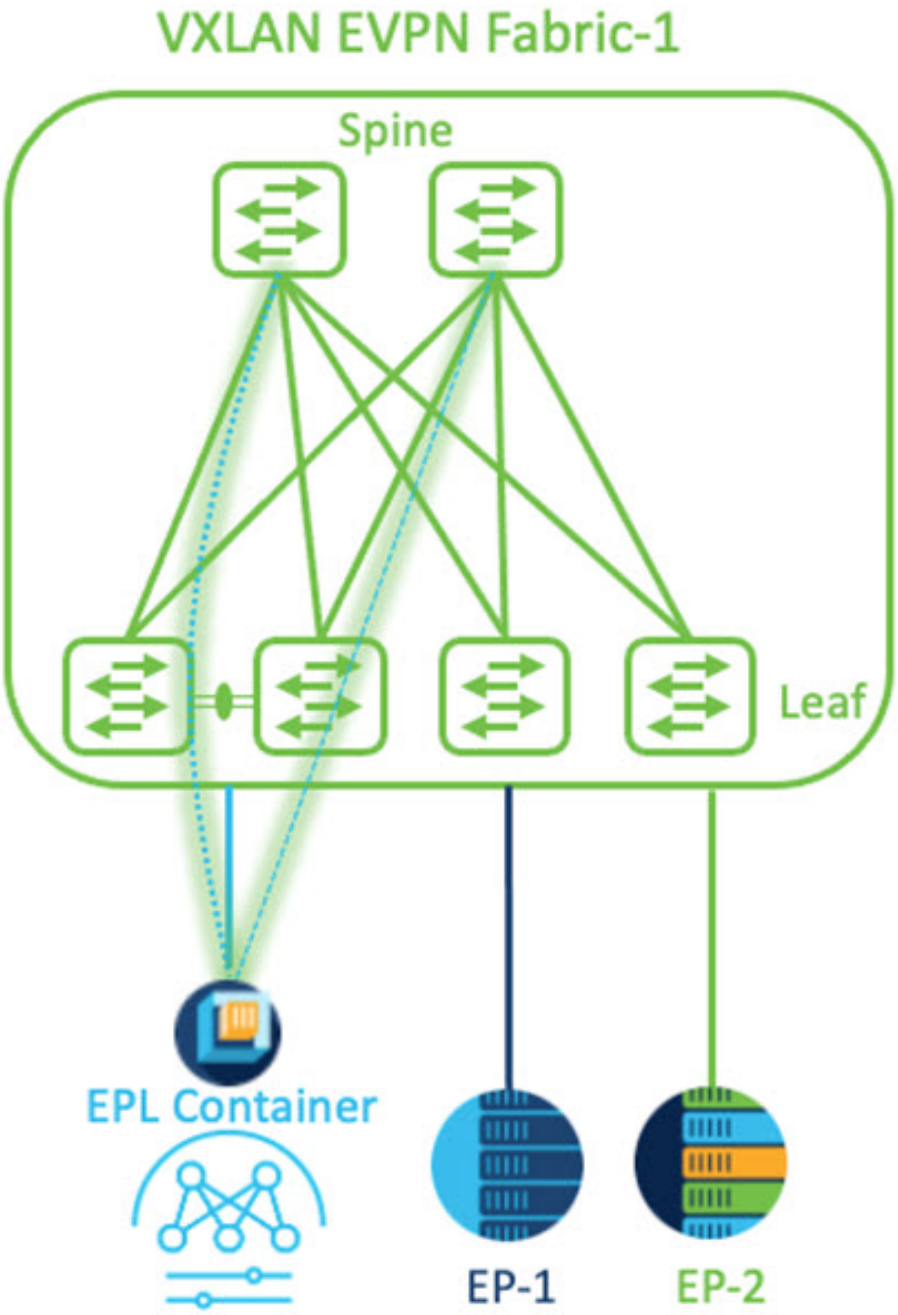
**Step 2** Click **Delete** to continue or **Cancel** to abort.

## Configuring Endpoint Locator for Single VXLAN EVPN Site

To configure endpoint locator for single VXLAN EVPN site, perform the following steps:

#### Before you begin

In the below figure, the NDFC service application is attached to the VPC pair of Leaf switches as it provides the link and node-level redundancy. The BGP instance running on EPL container establishes iBGP peering with the fabric spines. The iBGP peering is between Spine loopback addresses (loopback0) and EPL container persistent IP addresses. The loopback0 address of Spines is reachable via VXLAN Underlay, therefore, EPL container IP must have IP reachability towards the spines. We can configure an SVI on Leaf switches that can provide IP connectivity. The SVI will be a non-VXLAN enabled VLAN and will only participate in the underlay.




**Procedure**

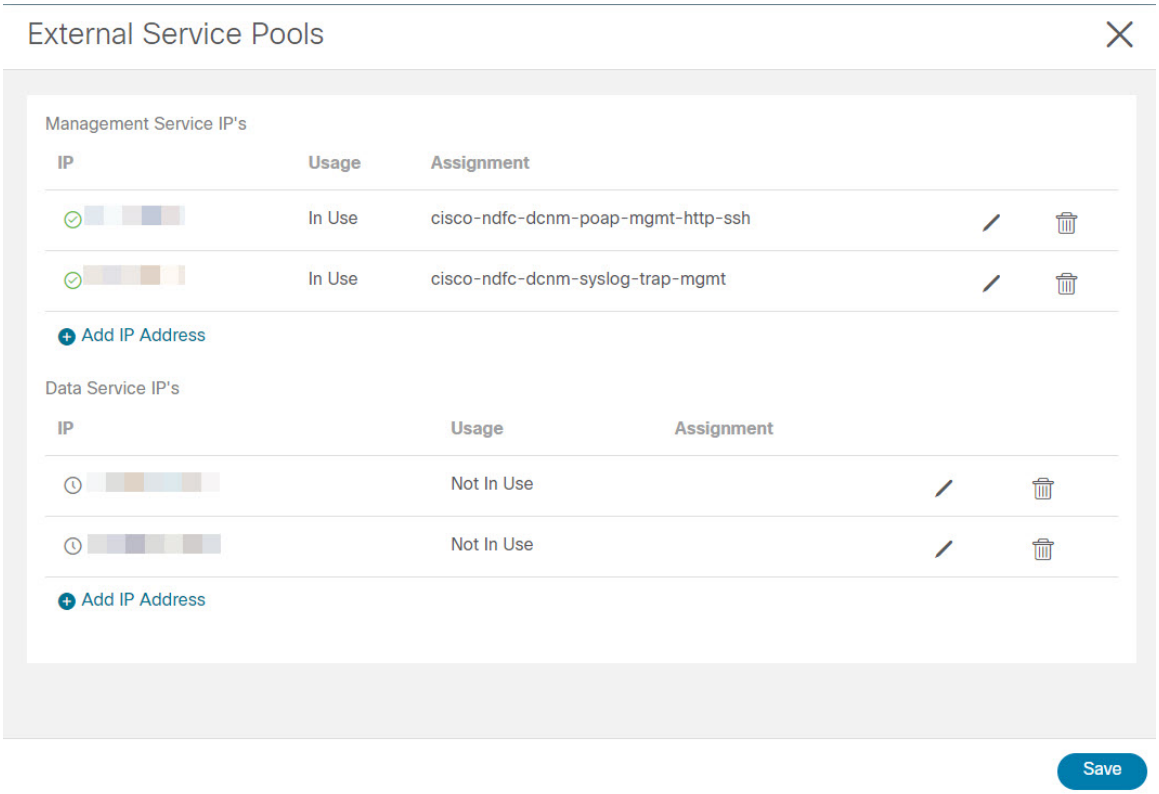
**Step 1** You must configure persistent IP addresses on Cisco Nexus Dashboard. On Nexus Dashboard, choose **Admin Console** > **Infrastructure** > **Cluster Configuration**.

**Step 2** On **General** tab, in **External Service Pools** card, click **Edit** icon.

The **External Service Pools** window appears.

**Step 3** Enter Persistent IP addresses in **Data Service IP's** and click **check** icon.

**Note**
The IP address must be associated with Nexus Dashboard Data Pool. A single persistent IP address is required to visualize and track EPs for a single site.




**Step 4** Configure SVI using FHRP for ND Data Interface and Underlay IP connectivity.

You can use **switch_freeform** policy on fabric Leaf 1.

To create a freeform policy, perform the following steps:

a) Choose **LAN**> **Fabrics**, double-click on required fabric.

The **Fabric Overview** page appears.

b) Click **Policy** tab, choose **Actions**> **Add Policy**.

The **Add Policy** window appears.

c) Choose appropriate Leaf1 switch from the **Switch List** drop-down list and click **Choose Template**.

d) On **Select Policy Template** window, choose **switch_freeform** template and click **Select**.

**Apply FHRP configurations and save the template.**

**Deploy the template configuration.**

In this example, SVI 100 with HSRP gateway created on fabric Leaf 1. Similarly, repeat the steps for fabric Leaf 2.

Below mentioned configuration example:

```
feature hsrp
vlan 100
name EPL-Inband
interface Vlan100
  no shutdown
  no ip redirects
  ip address 192.168.100.252/24
  no ipv6 redirects
  ip router ospf 100 area 0.0.0.0
  hsrp 100
    ip 192.168.100.254
```

**Step 5** Verify IP reachability between Nexus Dashboard Data Interface and fabric switches.

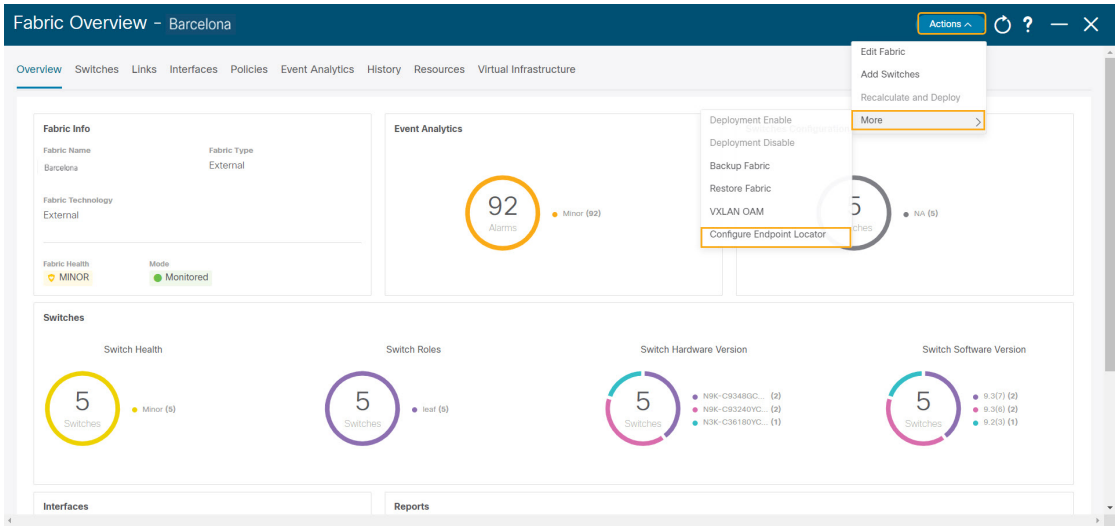
```
[rescue-user@ndfc-12-parth ~]$ ping 192.168.100.254 -c 2
PING 192.168.100.254 (192.168.100.254) 56(84) bytes of data.
64 bytes from 192.168.100.254: icmp_seq=1 ttl=255 time=1.95 ms
64 bytes from 192.168.100.254: icmp_seq=2 ttl=255 time=2.09 ms

--- 192.168.100.254 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 1.950/2.021/2.093/0.084 ms
[rescue-user@ndfc-12-parth ~]$
```

**Step 6** Enable EPL at fabric level.

a) To configure EPL, choose **LAN**> **Fabrics**> **Fabric Overview**.

b) On **Fabric Overview** window, choose **Actions**> **More**> **Configure EndPoint Locator**.




c) Choose the appropriate switches on the fabric hosting the Spine/Route Reflector RRs from the drop-down list.

Choose **Configure my Fabric** option for knob controls.

Whether BGP configuration will be pushed to the selected Spines/RRs as part of the enablement of the EPL feature. If the Spine/RR needs to be configured manually with a custom policy for the EPL BGP neighborship, then this option should be unchecked. For external fabrics that are only monitored and not configured on NDFC this option is grayed out. As these fabrics are not configured on NDFC.

Choose **Process MAC-Only Advertisements** option to enable processing of MAC-Only advertisements while configuring the EPL feature.

**Note**
If EPL is enabled on a fabric with or without choosing the **Process Mac-Only Advertisements** checkbox and if you want to toggle this selection later, then you must disable EPL and click **Database Clean-up** to delete endpoint data before re-enabling EPL with the desired **Process Mac-Only Advertisements** setting.

Choose **Yes** in **Collect Additional Information** to enable collection of additional information such as PORT, VLAN, and VRF while enabling the EPL feature. To access additional information, NX-API must be supported and enabled on the switches, ToRs, and leafs. If you choose **No** option, this information won't be collected and reported by EPL.

**Note**
For all fabrics except external fabrics, NX-API is enabled by default. For external fabrics, you must enable NX-API in the external fabric settings, choose **Enable NX-API** checkbox in the **Advanced** tab of the External_Fabric_11_1 fabric template.

Click on **Preview** icon to view a template of the configuration that is pushed to the switches enabling EPL. This configuration can be copied and pasted on spines or border gateway devices to enable EPL on external monitored fabrics.

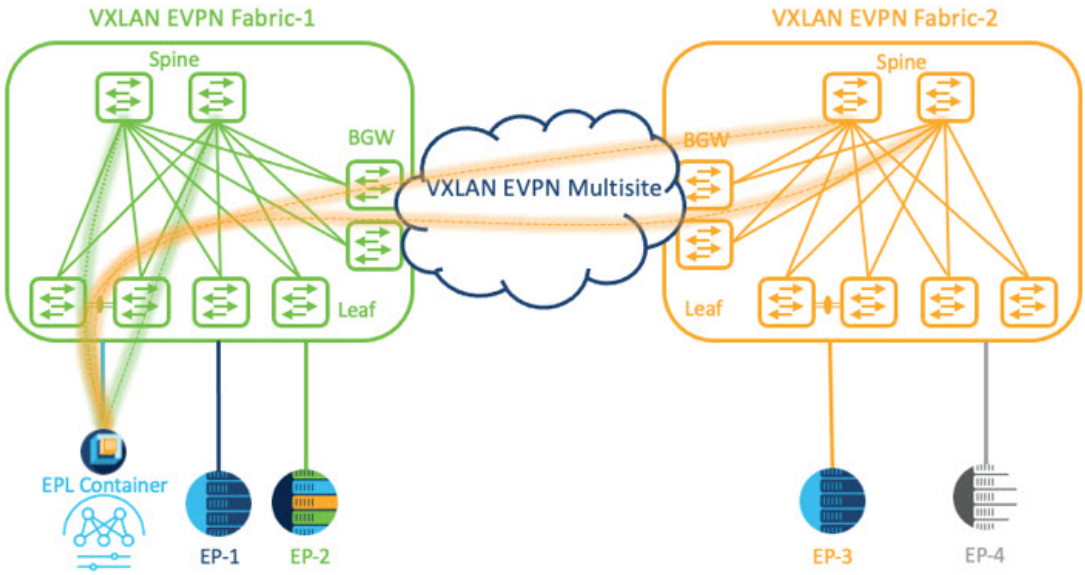
Once the appropriate selections are made and various inputs have been reviewed, click **Save Config** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message are displayed. Otherwise, EPL is successfully enabled. Once the EPL is enabled the Persistent IP will be in-use.

## Configuring Endpoint Locator for Multi-Fabric using VXLAN EVPN Multisite

To configure endpoint locator for mult-fabric VXLAN EVPN multisite, perform the following steps:

### Before you begin

The below figure enables EPL for Multi-Fabric using VXLAN EVPN Multisite. The BGP peering's are established between the Spines/RRs of each VXLAN EVPN Site and NDFC EPL Container. The Persistent IPs are required based on the number of VXLAN EVPN Sites. The NDFC application hosted on Cisco ND Cluster is located on Site 1. The routing information to reach the Spines/RRs deployed in the remote site must be exchanged across the Multisite. Once the BGP session is formed, local EPs of Fabric 2 can be visualized and tracked.




By default, Nexus Dashboard data Interface and Site 2 Spines/RRs loopback prefixes are not advertised across the BGWs. Therefore, prefixes must be exchanged using custom route maps and prefix lists across the sites. At the same time, route redistribution between OSPF and BGP is required as Spines/RRs loopback prefixes are part of OSPF protocol while BGWs peer with each other using BGP.

### Procedure

**Step 1** You must configure persistent IP addresses on Cisco Nexus Dashboard. On Nexus Dashboard, choose **Admin Console** > **Infrastructure** > **Cluster Configuration**.

**Step 2** On **General** tab, in **External Service Pools** card, click **Edit** icon.

The **External Service Pools** window appears.

**Step 3** Enter Persistent IP addresses in **Data Service IP's** and click **check** icon.

**Note**
Ensure that the IP addresses are associated with Nexus Dashboard Data Pool. Two persistent IP addresses are required to visualize and track EPs for a multisite with two member fabrics. One Persistent Data IP address is used as EPL container IP to establish BGP session with Site 1 fabric. A new Persistent IP address is configured that can be used to peer with Site 2 fabric.

**Step 4** Configure Route Redistribution for VXLAN EVPN Fabrics.

Route Redistribution for Fabric 1

The following switch_freeform policy can be used on Fabric 1 BGWs. To create a new **switch_freeform** policy, refer to the above examples.

Below the example of sample configuration

```
ip prefix-list site-2-rr seq 5 permit 20.2.0.1/32 >> Site 2 RR
ip prefix-list site-2-rr seq 6 permit 20.2.0.2/32 >> Site 2 RR
```

```
ip prefix-list epl-subnet seq 5 permit 192.168.100.0/24 >> EPL Subnet

route-map bgp-to-ospf permit 10
   match ip address prefix-list site-2-rr
route-map ospf-to-bgp permit 10
   match ip address prefix-list epl-subnet

router ospf 100
   redistribute bgp 100 route-map bgp-to-ospf

router bgp 100
   address-family ipv4 unicast
       redistribute ospf 100 route-map ospf-to-bgp
```

Route Redistribution for Fabric 2

The following switch_freeform policy can be used on Fabric 2 BGWs. To create a new **switch_freeform** policy, refer to the above examples.

Below the example of sample configuration

```
ip prefix-list site-2-rr seq 5 permit 20.2.0.1/32 >> Site 2 RR
ip prefix-list site-2-rr seq 6 permit 20.2.0.2/32 >> Site 2 RR
ip prefix-list epl-subnet seq 5 permit 192.168.100.0/24 >> EPL Subnet

route-map bgp-to-ospf permit 10
  match ip address prefix-list epl-subnet
route-map ospf-to-bgp permit 10
  match ip address prefix-list site-2-rr

router ospf 200
  redistribute bgp 200 route-map bgp-to-ospf

router bgp 200
    address-family ipv4 unicast
        redistribute ospf 200 route-map ospf-to-bgp
```

**Step 5** To configure EPL, choose **LAN**> **Fabrics**> **Fabric Overview**.

**Step 6** On **Fabric Overview** window, choose **Actions**> **More**> **Configure EndPoint Locator**.

**Step 7** Choose the appropriate switches on the fabric hosting the Spine/Route Reflector RRs from the drop-down list.

Once the appropriate selections are made and various inputs have been reviewed, click **Save Config** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled. Once the EPL is enabled the Persistent IP will be in-use.

You can view EPL enabled for fabric-1 and fabric-2 successfully. To view and track EPs, Refer the Monitoring Endpoint Locator section.

## Configuring Endpoint Locator for vPC Fabric Peering Switches

Networks Administrator can create vPC between a pair of switches using a Physical Peer Link or Virtual Peer link. vPC Fabric Peering provides an enhanced dual-homing access solution without the overhead of wasting physical ports for vPC Peer Link. For Virtual Peer link, EPL can still be connected to vPC pair of Leaf switches for the link and node-level redundancy. However, VXLAN VLAN (Anycast Gateway) as the First hop for EPL will be used. The loopback0 address of Spines/RRs is reachable only via VXLAN Underlay, while VXLAN VLAN will be part of a Tenant VRF. Therefore, to establish IP communication, route-leaking is configured between Tenant VRF and Default VRF. For more information, refer to vPC Fabric Peering section.

To configure endpoint locator for vPC Fabric Peering switches perform the following steps:

**Procedure**

**Step 1** You must configure persistent IP addresses on Cisco Nexus Dashboard. On Nexus Dashboard, choose **Admin Console** > **Infrastructure** > **Cluster Configuration**.

**Step 2** On **General** tab, in **External Service Pools** card, click **Edit** icon.

The **External Service Pools** window appears.

**Step 3** Enter Persistent IP addresses in **Data Service IP's** and click **check** icon.

**Step 4** Create a Tenant VRF and Anycast Gateway on the vPC fabric peering switches.

add two images

**Step 5** Configure Route-leaking between Tenant VRF and Default VRF.

**Advertise from Tenant VRF to Default VRF.**

The following switch_freeform policy can be used on fabric Leaf where ND is connected.

```
ip prefix-list vrf-to-default seq 5 permit 192.168.100.0/24 >> EPL subnet
route-map vrf-to-default permit 10
  match ip address prefix-list vrf-to-default
vrf context epl_inband
  address-family ipv4 unicast
    export vrf default map vrf-to-default allow-vpn
router ospf UNDERLAY
   redistribute bgp 200 route-map vrf-to-default
```

**Advertise from Default VRF to Tenant VRF.**

The following switch_freeform policy can be used on fabric Leaf where ND is connected.

```
ip prefix-list default-to-vrf seq 5 permit 20.2.0.3/32 >> Spine loopback IP
ip prefix-list default-to-vrf seq 6 permit 20.2.0.4/32 >> Spine loopback IP
route-map default-to-vrf permit 10
   match ip address prefix-list default-to-vrf
vrf context epl_inband
   address-family ipv4 unicast
       import vrf default map default-to-vrf
        router bgp 200
   address-family ipv4 unicast
       redistribute ospf UNDERLAY route-map default-to-vrf
```

**Step 6** Enable EPL at fabric level.

a) To configure EPL, choose **LAN**> **Fabrics**> **Fabric Overview**.

b) On **Fabric Overview** window, choose **Actions**> **More**> **Configure EndPoint Locator**.

c) Choose the appropriate switches on the fabric hosting the Spine/Route Reflector RRs from the drop-down list.

Once the appropriate selections are made and various inputs have been reviewed, click **Save Config** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled. Once the EPL is enabled the Persistent IP will be in-use.

# Configuring Endpoint Locator for External Fabrics

In addition to Easy fabrics, Nexus Dashboard Fabric Controller allows you to enable EPL for VXLAN EVPN fabrics comprising of switches that are imported into the external fabric. The external fabric can be in managed mode or monitored mode, based on the selection of **Fabric Monitor Mode** flag in the **External Fabric** Settings. For external fabrics that are only monitored and not configured by Nexus Dashboard Fabric Controller, this flag is disabled. Therefore, you must configure BGP sessions on the Spine(s) via OOB or using the CLI. To check the sample template, click  icon to view the configurations required while enabling EPL.

In case the **Fabric Monitor Mode** checkbox in the External Fabric settings is unchecked, then EPL can still configure the spines/RRs with the default **Configure my fabric** option. However, disabling EPL would wipe out the router bgp config block on the spines/RRs. To prevent this, the BGP policies must be manually created and pushed onto the selected spines/RRs.

# Configuring Endpoint Locator for eBGP EVPN Fabrics

You can enable EPL for VXLAN EVPN fabrics, where eBGP is employed as the underlay routing protocol. Note that with an eBGP EVPN fabric deployment, there is no traditional RR similar to iBGP. The reachability of the in-band subnet must be advertised to the spines that behave as Route Servers. To configure EPL for eBGP EVPN fabrics from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN > Fabrics**.

Select the fabric to configure eBGP on or create eBGP fabric with the **Easy_Fabric_eBGP** template.

**Step 2** Use the **leaf_bgp_asn** policy to configure unique ASNs on all leaves.

**Step 3** Add the **ebgp_overlay_leaf_all_neighbor** policy to each leaf.

Fill **Spine IP List** with the spines' BGP interface IP addresses, typically the loopback0 IP addresses.

Fill **BGP Update-Source Interface** with the leaf's BGP interface, typically loopback0.

**Step 4** Add the **ebgp_overlay_spine_all_neighbor** policy to each spine.

Fill **Leaf IP List** with the leaves' BGP interface IPs, typically the loopback0 IPs.

Fill **Leaf BGP ASN** with the leaves' ASNs in the same order as in **Leaf IP List**.

Fill **BGP Update-Source Interface** with the spine's BGP interface, typically loopback0.

After the in-band connectivity is established, the enablement of the EPL feature remains identical to what is listed so far. EPL becomes a iBGP neighbor to the Route Servers running on the spines.

# Monitoring Endpoint Locator

Information about the Endpoint Locator is displayed on a single landing page or dashboard. The dashboard displays an almost real-time view of data (refreshed every 30 seconds) pertaining to all the active endpoints on a single pane. The data that is displayed on this dashboard depends on the scope selected by you from the

**SCOPE** drop-down list. The Nexus Dashboard Fabric Controller scope hierarchy starts with the fabrics. Fabrics can be grouped into a Multi-Site Domain (MSD). A group of MSDs constitute a Data Center. The data that is displayed on the Endpoint Locator dashboard is aggregated based on the selected scope. From this dashboard, you can access Endpoint History, Endpoint Search, and Endpoint Life.

**Note** This is a preview feature in Nexus Dashboard Fabric Controller, Release 12.0.1a. We recommend that you use this feature marked as BETA in your lab setup only. Do not use these features in your production deployment.

# Disabling Endpoint Locator

To disable endpoint locator from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **Endpoint Locator > Configure**.

The **Endpoint Locator** window appears. Select the required fabric from the **SCOPE** dropdown list. The fabric configuration details are then displayed for the selected fabric.

**Step 2** Click **Disable**.

# Fabric Overview

The **Actions** drop-down list at the Fabric level allows you to perform the following:

| Actions | Description |
|---|---|
| Edit Fabric | • To edit a fabric, choose **Actions** > **Edit Fabric**.<br>• The Edit fabric window appears, do necessary changes and click **Save**. |
| Add Switches | Refer to section Add Switches for more information. |
| Recalculate Config | Refer to section Recalculate Config for more information. |
| Preview Config | Refer to section Preview Config for more information. |
| Deploy Config | • To deploy configuration changes, choose **Actions** > **Deploy Config**.<br>• A progress window appears and confirmation message is displayed. |
| **More** | |

| Actions | Description |
|---|---|
| Deployment Enable | • From Fabrics Overview, choose **Actions** on main tab, choose **More** > **Deployment Enable**.<br>• A confirmation window appears, click **OK**. |
| Deployment Disable | • From Fabrics Overview, choose **Actions** on main tab, choose **More** > **Deployment Disable**.<br>• A confirmation window appears, click **OK**. |
| Backup Fabric | Refer to Backup Fabric section for more information. |
| Restore Fabric | Refer to Restore Fabric section for more information. |
| VXLAN OAM | Refer to VXLAN OAM, on page 164 section for more information.<br>**Note**<br>This feature appears in the Actions drop-down list only for VXLAN Fabric, eBGP VXLAN Fabric, External, and Lan Classic fabric technologies, which support VXLAN OAM. |
| Configure End Point Locator | The Endpoint Locator (EPL) feature allows real-time tracking of endpoints within a data center. For more information, see Endpoint Locator , on page 166.<br>**Note**<br>This is a preview feature in Nexus Dashboard Fabric Controller, Release 12.0.1a. We recommend that you use this feature marked as BETA in your lab setup only. Do not use these features in your production deployment. |

Fabric Overview contains tabs that allows you view and perform all the operations on the fabric.

# Overview

The **Overview** tab displays the following information as cards.

- Fabric Information
- Fabrics

  Displayed if there are child fabrics. For example: Multi-Site Fabrics
- Event Analytics
- Switches Configuration
- Switches
  - Switch Health
  - Switch Configuration
  - Switch Roles

    - Switch Hardware Version

- VXLAN

  Displayed only for VXLAN Fabrics

    - Routing Loopback
    - VTEP Loopback
    - Multisite Loopback
    - NVE Int Status
    - Networks/VRFs Definition
    - Extended Networks/VRFs

- Hosts

  This tab is displayed only in you've configured IPFM fabric.

- Flows

  This tab is displayed only in you've configured IPFM fabric.

- Reports

### Hosts

The **Hosts** card displays the following details:

- **Pie chart** - Each slice has a unique color and displays a host role and count, for example, Sender, Receiver, and ARP. Click a host type, for example, Sender, to hide or unhide the slice, for the selected IPFM fabric.

  To view more information, choose **Fabric Overview** > **Hosts** > **Discovered Hosts**.

- **Faults** - If faults exist, displays the number of faults including policer drops. To view more information, click **Faults** which opens the **Hosts** > **Discovered Hosts** tab.

For more information about hosts, see Hosts, on page 224.

### Flows

The **Flows** card displays the following details:

- **Pie chart** - Each slice has a unique color and displays a multicast flow class and count, for example, Active, Inactive, Sender Only, and Receiver Only. Click a flow class, for example, Active, to hide or unhide the slice.

  To view more information, choose **Fabric Overview** > **Flows** > **Flow Status**.

- **Groups** - Displays the number of multicast flow groups. This information is also displayed on the IPFM fabric topology.

For more information about flows, see Flows, on page 235.

# Switches

You can manage switch operations in this tab. Each row represents a switch in the fabric, and displays switch details, including its serial number.

Some of the actions that you can perform from this tab are also available when you right-click a switch in the fabric topology window. However, the **Switches** tab enables you to provision configurations on multiple switches, like deploying policies, simultaneously.

The Switches tab has following information of every switch you discover in the fabric:

- Name: Specifies the switch name.
- IP Address: Specifies the IP address of the switch.
- Role: Specifies the role of the switch.
- Serial Number: Specifies the serial number of the switch.
- Fabric Name: Specifies the name of the fabric, where the switch is discovered.
- Fabric Status: Specifies the status of the fabric, where the switch is discovered.
- Discover Status: Specifies the discovery status of the switch.
- Model: Specifies the switch model.
- Software Version: Specifies the software version of the switch.
- Last Updated: Specifies when the switch was last updated.
- Mode: Specifies the current mode of the switch.
- vPC Role: Specifies the vPC role of the switch.
- vPC Peer: Specifies the vPC peer of the switch.

The **Switches** tab has the following operations on the Action drop-down list:

- **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.

- Preview: You can preview the pending configurations and the side-by-side comparison of running configurations and expected configurations.
- Deploy: Deploy switch configurations. From Cisco Nexus Dashboard Fabric Controller 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.

- Discovery: You can perform the following operations.
  - Update discovery credentials: Update device credentials such as authentication protocol, username and password.
  - Rediscover switch: Initiate the switch discovery process by Nexus Dashboard Fabric Controller afresh.
- 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 in the External Fabric: 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.



**Note** Note: NDFC 12 does not allow you to create vPC pairing on Spine, Border Spine, Border Gateway Spine, Super Spine, Border Super Spine, and Border Gateway Super Spine roles.

- vPC Overview
- More: Further operations are provided under More.
- Show Commands: Execute Show commands on the selected Switch. Select the Commands from the drop-down list. Enter appropriate values in the Variables fields, and click **Execute**. The right column execute the show command and displays the output.

- Exec Commands: 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.
- Provision RMA: Allows you to replace a physical switch in a Fabric when using Cisco Nexus Dashboard Fabric Controller Easy Fabric mode.
- Copy Run Start: 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.

- Reload: Reload the selected switch.



**Note** This option is grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.

- Delete switches: Remove the switch from the fabric.

  This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.
- Restore Switches: 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.
- Change Mode: You can change the mode of the switch from Normal to Managed and vice versa.

  You can choose to save the settings and deploy immediately, or schedule it for later.

## Guidelines and Limitations for Changing Discovery IP Address

From Cisco Nexus Dashboard Fabric Controller Release 12.0.1a, you can change the Discovery IP address of a device that is existing in a fabric.

### Guidelines and Limitations

The following are the guidelines and limitations for changing discovery IP address.

- Changing discovery IP address is supported for NX-OS switches and devices that are discovered over their management interface.
- Changing discovery IP address is supported for templates such as:
  - Easy_Fabric
  - Easy_Fabric_eBGP

- External
- LAN_Classic
- LAN_Monitor

- Changing discovery IP address is supported in both managed and monitored modes.
- Only users with the **network-admin** role can change the discovery IP address on Cisco Fabric Controller UI.
- The discovery IP address must not be used on other devices, and it must be reachable when the change is done.
- While changing the discovery IP address for a device in a managed fabric, switches are placed in migration mode.
- When you change the IP address of a switch that is linked to vPC Peer, corresponding changes such as vPC peer, domain configuration will be updated accordingly.
- Fabric configuration restores the original IP address, it reports out of sync post restore and the configuration intent for the device must be updated manually to get the in-sync status.
- Fabric controllers restore that had the original device discovery IP reports the switch as Unreachable post restore. The discovery IP address change procedure must be repeated after the restore.
- Device Alarms associated with the original discovery IP address will be purged after the change of IP address.

## Changing Discovery IP Address

### Before you begin

You must make the management IP address and route related changes on the device and ensure that the reachability of the device from Nexus Dashboard Fabric Controller.

To change the discovery IP address from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Procedure

**Step 1** Choose **LAN** > **Fabrics**.

**Step 2** Click on fabric names to view the required switch.

The **Fabric summary** slide-in pane appears.

**Step 3** Click **Launch** icon to view **Fabric Overview** window.

**Step 4** On the **Switches** tab, click **Refresh** icon adjacent to the **Action** button on the main window.

Switch with a changed IP address will be in **Unreachable** state in **Discovery Status** column.

**Step 5** Click the check box next to the **Switch** column and select the switch.

**Note**

You can change the IP address for individual switch and not for multiple switches.

**Step 6** Choose **Actions** > **Change Discovery IP** on the switches tab area.

The **Change Discovery IP** window appears.

Similarly, you can navigate from **LAN** > **Switches** tab. Choose a required switch, click **Actions** > **Discovery** > **Change Discovery IP**.

**Step 7** Enter the appropriate IP address in the **New IP Address** text field and click **OK**.

a) The new IP address must be reachable from Nexus Dashboard Fabric Controller to update successfully.
b) Repeat the above procedures for the devices where the discovery IP address must be changed before proceeding with further steps.
c) If the fabric is in managed mode, the device mode will be updated to migration mode.

**Step 8** From the fabric **Actions** drop-down list, click **Recalculate Config** to initiate the process of updating Nexus Dashboard Fabric Controller configuration intent for the devices. Similarly, you can recalculate configuration on topology window. Choose **Topology**, tab right-click on the switch, click **Recalculate Config**.

The Nexus Dashboard Fabric Controller configuration intent for the device management related configuration will be updated and the device mode status for the switch is changed to normal mode. The switch configuration status is displayed as **In-Sync**.

**Note**
The PM records associated with the old switch IP address will be purged and new record collections take an hour to initiate after the changes.

# 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 Nexus Dashboard Fabric Controller.

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 color 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**.

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.

Starting from Cisco NDFC Release 12.1.2e, parameters MTU, SPEED, Source Interface Description, Destination Interface Description, Source Interface Freeform Config, and Destination Interface Freeform Config are added to the existing **int_pre_provision_intra_fabric_link** template. These parameters are preserved on subsequent **Recalculate & Deploy** after the device has completed bootstrap and POAP.

The following table describes the fields that appear on **Links** tab.

| Field | Description |
|---|---|
| Fabric Name | Specifies the name of the Fabric. |
| Name | Specifies the name of the link.<br>The list of previously created links is displayed. The list contains intra-fabric links, which are between switches within a fabric, and inter-fabric links, which are between border switches in this fabric and switches in other fabrics. |
| Policy | Specifies the link policy. |
| Info | Provides more information about the link. |
| Admin State | Displays the administrative state of the link. |
| Oper State | Displays the operational state of the link. |

The following table describes the action items, in the Actions menu drop-down list, that appear on **Fabric Overview > Links > Links**.

| Action Item | Description |
|---|---|
| Create | Allows you to create the following links:<br>• Creating Inter-Fabric Links, on page 191<br>• Creating Intra-Fabric Links, on page 190 |
| Edit | Allows you to edit the selected fabric. |
| Delete | Allows you to delete the selected fabric. |
| Import | 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.<br>**Note**<br>• You cannot update existing links.<br>• The Import Links icon is disabled for external fabric. |
| Export | Choose the link and select **Export** to export the links in a CSV file.<br>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. |

# Creating Intra-Fabric Links

Click the Links tab. You can see a list of links. The list is empty when you are yet to create a link.

To create Intra-Fabric links, perform the following steps:

### Procedure

**Step 1** From the Actions drop-down list, select **Create**.

The **Link Management - Create Link** page appears.

**Step 2** From the Link Type drop-down box, choose **Intra-Fabric** since you are creating an IFC. The screen changes correspondingly.

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**: If the link is between two ethernet interfaces assigned with IP addresses, choose int_intra_fabric_num_link.
- **int_intra_fabric_unnum_link**: If the link is between two IP unnumbered interfaces, choose int_intra_fabric_unnum_link.
- **int_intra_vpc_peer_keep_alive_link**: If the link is a vPC peer keep-alive link, choose int_intra_vpc_peer_keep_alive_link.
- **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 sate of the interface.

MTU – Specify the maximum transmission unit (MTU) through the two interfaces.

Source Interface Description and Destination Interface Description – Describe the links for later use. For example, if the link is between a leaf switch and a route reflector device, you can enter the information in these fields (Link from leaf switch to RR 1 and Link from RR 1 to leaf switch). This description will converted into a config, but will not be pushed into the switch. After Save & Deploy, it will reflect in the running configuration.

**Disable BFD Echo on Source Interface** and **Disable BFD Echo on Destination Interface** – Select the check box to disable BFD echo packets on source and destination interface.

Note that the BFD echo fields are applicable only when you have enabled BFD in the fabric settings.

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

**Step 3** Click **Save** at the bottom right part of the screen.

You can see that the IFC is created and displayed in the list of links.

**Step 4** On the Fabric Overview Actions drop-down list, select **Recalculate Config**.

The Deploy Configuration 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 **Pending Config** column. It lists the pending configurations on the switch. The Side-by-side Comparison tab displays the running configuration and expected configuration side-by-side.

Close the **Pending Config** screen.

**Step 5** From **Fabric Overview Actions** drop-down list, select **Deploy Config**.

The pending configurations are deployed.

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. In the fabric topology, you can see that the link between the two devices is displayed.

## Creating Inter-Fabric Links

Click the Links tab. You can see a list of links. The list is empty when you are yet to create a link.



**Note** In external fabrics, inter-fabric links support BGW, Border Leaf/Spine, and edge router switches.

To create Inter-Fabric links, perform the following steps:

### Procedure

**Step 1** From the Actions drop-down list, select **Create**.

The **Link Management - Create Link** page appears.

**Step 2** From the Link Type drop-down box, choose **Inter-Fabric** since you are creating an IFC. The screen changes correspondingly.

The fields for inter-fabric link creation are as follows:

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

Link Sub-Type – This field populates the IFC type. From the drop-down list, choose **VRF_LITE**, **MULTISITE_UNDERLAY**, or **MULTISITE_OVERLAY**.

The Multi-Site options are explained in the Multi-Site use case.

For information about VXLAN MPLS interconnection, see the MPLS SR and LDP Handoff, on page 619 chapter.

For information about routed fabric interconnection, see *Creating Inter-Fabric Links Between a Routed Fabric and an External Fabric* section in *Configuring a Fabric with eBGP Underlay* chapter.

**Link Template** – The link template is populated.

The templates are autopopulated with corresponding prepackaged default templates that are based on your selection.

**Note**
You can add, edit, or delete user-defined templates. See Templates 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 that is 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.

**Step 3** Click **Save** at the bottom-right part of the screen.

You can see that the IFC is created and displayed in the list of links.

**Step 4** On the Fabric Overview Actions drop-down list, select **Recalculate Config**.

The Deploy Configuration 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 **Pending Config** column. It lists the pending configurations on the switch. The Side-by-side Comparison tab displays the running configuration and expected configuration side by side.

Close the **Pending Config** screen.

**Step 5** From the **Fabric Overview Actions** drop-down list, select **Deploy Config**.

The pending configurations are deployed.

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. 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.

### What to do next

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 over MSD fabric, IFCs are automatically created between the (VXLAN fabric) border/BGW device and connected (external fabric) Edge router or Core device. When you remove the ER/core/border/BGW device, the corresponding IFCs (link PTIs) to/from that switch are deleted on Nexus Dashboard Fabric Controller. Next, Nexus Dashboard Fabric Controller 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, select the VXLAN fabric and the extended VRFs, and undeploy the VRFs in the VRF deployment screen.

To delete the IFCs, 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 Nexus Dashboard Fabric Controller, the underlay networks that are provisioned on those switches, and the configurations between Nexus Dashboard Fabric Controller 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.

# Interfaces

This section contains the following topics:


- Interfaces, on page 317
- Interface Groups, on page 330


# Policies

Nexus Dashboard Fabric Controller provides the ability to group a set of switches, and allows you to push a set of underlay configurations to the group.

Choose **LAN > Policies** to display the list of policies.

The following table describes the fields that appear on **Fabric Overview > Policies**.

| Field | Description |
|---|---|
| Policy ID | Specifies the policy ID. |
| Switch | Specifies the switch name. |
| IP Address | Specifies the IP address of the switch. |
| Template | Specifies the name of the template. |
| Description | Specifies the description. |
| Entity Name | Specifies the entity name. |
| Entity Type | Specifies the entity type. |
| Source | Specifies the source. |
| Priority | Specifies the priority. |
| Content Type | Species for the content type. |
| Fabric Name | Specifies the fabric name. |
| Serial Number | Specifies the serial number of the switch. |
| Editable | Specifies a Boolean value to indicate if the policy is editable. |
| Mark Deleted | Specifies a Boolean value to indicate if the policy is marked to be deleted. |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **LAN > Policies**.

| Action Item | Description |
|---|---|
| Add Policy | |

| Action Item | Description |
|---|---|
| Edit Policy | Choose a policy from the table and choose **Edit Policy** to modify the policy.<br>**Note**<br>• The policies in the italics font cannot be edited. The value under the **Editable** and **Mark Deleted** columns for these policies is false.<br>• A warning appears when you edit a policy whose **Mark Deleted** value is set to *true*. The switch freeform child policies of **Mark Deleted** policies appears in the Policies dialog box. You can edit only **Python** switch_freeform policies. You cannot edit **Template_CLI** switch_freeform_config policies. |
| Delete Policy | Choose policies from the table and choose **Delete Policy** to delete the policies.<br>**Note**<br>A warning appears when you delete policies whose **Mark Deleted** values are set to *true*.<br>When you delete the policies whose **Mark Deleted** values are set to *true*, these entries are deleted from the NDFC database only but the configs are not deployed to the switch. |
| Generated Config | Choose policies from the table and choose **Generated Config** to view the delta of configuration changes made by every user. |
| Push Config | Choose policies from the table and choose **Push Config** to push the policy configuration to the device.<br>**Note**<br>• This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.<br>• A warning appears if you push configuration for a Python policy.<br>• A warning appears when you push configurations for policies whose **Mark Deleted** values are set to *true*. |

## Adding a Policy

To add a policy, perform the following steps:

### Procedure

**Step 1** Choose **Actions** > **Add Policy**.

The **Create Policy** window appears.

**Step 2** Click and choose required switch and click **Select**.

You must deploy the switch in a pending state.

**Step 3** Click **Choose Template** and choose appropriate policy template and click **Select**.

From Cisco NDFC Release 12.1.2e, you can enable or disable PTP high-correction notification when the system encounters a high-correction event. Whenever the correction value exceeds the configured value then that correction is called a high-correction. By default, a high-correction notification is disabled. Enable it manually to generate the notification. Perform the following steps to enable the high-correction notification:

a. **Enable PTP Telemetry** – Check this check box to enable telemetry for PTP.

b. **Is Large-Scale Fabric?** – Check this check box to generate the high-correction notification. Are there more than 35 devices in the fabric. If yes, PTP events will be used if the switch version is 9.3(5) or higher, or else PTP correction data will be pushed periodically.

c. **PTP High-Correction Interval** – Specify the wait time between two successive notifications, duration value is in seconds.

d. **PTP Correction Range** – Set correction range threshold value (ns), default is 100000 (100us).

From Cisco NDFC Release 12.1.2e, new templates **ipv4_prefix_list** and **ipv6_prefix_list** are added to the template list.

**Step 4** Enter the required name in the **Prefix List Name** field. Perform the following steps to include the prefix-list entries:

a. On the **Prefix-list Entries** field, click **Actions** > **Add**.

The **Add Item** window appears.

b. The mandatory fields on the **Add Item** window are:

**IPv4 Prefix** – Enter the ipv4 prefix address.

**Sequence Number** – Enter the value in the sequence number.

**Action** – From the drop-down list, choose **permit** or **deny**.

Click **Save**.

**Step 5** Repeat the step (5) to add the required number of prefix-list entries.

**Note**
The value in the **Sequence Number** must be higher than the previous prefix-list entry. If not, an error message is displayed.

**Step 6** (Optional) Select the required prefix-list entry and click **Actions > Edit** to edit the selected prefix-list entry.

**Step 7** (Optional) Select the appropriate prefix-list entry and click **Actions > Insert Above** to insert a new prefix-list entry.

**Note**
The value in the **Sequence Number** must be lower than the below prefix-list entry. If not, an error message is displayed.

**Step 8** 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.

**Step 9** Enter the mandatory parameters in the text field and click **Save**.

## Advertising PIP on vPC

Choose required fabric on LAN Fabric window and Navigate to **Edit Fabric** > **VPC**, check the **vPC advertise-pip** check box to enable the Advertise PIP feature on all vPCs in a fabric. Choose 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.
- The policy vpc_advertise_pip_jython can be applied only when switches are part of vPC pairing.
- Ensure that you configure **vpc advertise-pip** command during maintenance window as it involves BGP next-hop rewrite. Enabling this feature EVPN type 5 uses Switch Primary IP as next-hop while EVPN type 2 continue to use Secondary IP.
- 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** Choose **LAN** > **Fabrics**. Double-click on the required fabric.

The **Fabric Overview** window appears.

**Step 2** On the **Fabric Overview** window, choose **Policies** > **Add Policy** and then select a switch with vPC.

**Step 3** Click **Actions > Add** and choose the switch from the **Switch List** drop-down list. Choose **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 4** Click **Save**, and then deploy this policy.

## Viewing and Editing Policies

Nexus Dashboard Fabric Controller provides the ability to group a set of switches, and allows you to push a set of underlay configurations to the group.

Choose **LAN > Policies** to display the list of policies.

The following table describes the fields that appear on **LAN > Policies**.

| Field | Description |
|---|---|
| Policy ID | Specifies the policy ID. |
| Switch | Specifies the switch name. |
| IP Address | Specifies the IP address of the switch. |
| Template | Specifies the name of the template. |
| Description | Specifies the description. |
| Entity Name | Specifies the entity name. |
| Entity Type | Specifies the entity type. |
| Source | Specifies the source. |
| Priority | Specifies the priority. |
| Content Type | Species for the content type. |
| Fabric Name | Specifies the fabric name. |
| Serial Number | Specifies the serial number of the switch. |
| Editable | Specifies a Boolean value to indicate if the policy is editable. |
| Mark Deleted | Specifies a Boolean value to indicate if the policy is marked to be deleted. |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **LAN > Policies**.

| Action Item | Description |
|---|---|
| Add Policy | To add a policy, see Add Policy |

| Action Item | Description |
|---|---|
| Edit Policy | Choose a policy from the table and choose **Edit Policy** to modify the policy.<br>**Note**<br>• The policies in the italics font cannot be edited. The value under the **Editable** and **Mark Deleted** columns for these policies is false.<br>• A warning appears when you edit a policy whose **Mark Deleted** value is set to *true*. The switch freeform child policies of **Mark Deleted** policies appears in the Policies dialog box. You can edit only **Python** switch_freeform policies. You cannot edit **Template_CLI** switch_freeform_config policies. |
| Delete Policy | Choose policies from the table and choose **Delete Policy** to delete the policies.<br>**Note**<br>A warning appears when you delete policies whose **Mark Deleted** values are set to *true*. |
| Generated Config | Choose policies from the table and choose **Generated Config** to view the delta of configuration changes made by every user. |
| Push Config | Choose policies from the table and choose **Push Config** to push the policy configuration to the device.<br>**Note**<br>• This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.<br>• A warning appears if you push configuration for a Python policy.<br>• A warning appears when you push configurations for policies whose **Mark Deleted** values are set to *true*. |

## Custom Maintenance Mode Profile Policy

When you place a switch in maintenance mode using NDFC, only a fixed set of BGP and OSPF isolate CLIs are configured in the maintenance mode profile. 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 Custom Maintenance Mode Profile Policy

To create and deploy a custom maintenance mode profile policy from **Web UI > Switches**, perform the following procedure.

### Procedure

**Step 1** Select the desired switch and launch **Switch Overview**.

**Step 2** On the Policies tab, select **Actions > Add Policy** to add a new policy.

**Step 3** On the Create Policy screen, click **Choose Template**.

**Step 4** Select **custom_maintenance_mode_profile** from the **Select Policy Template** list.

**Step 5** Fill in the **Maintenance mode profile contents** with the desired configuration CLIs.

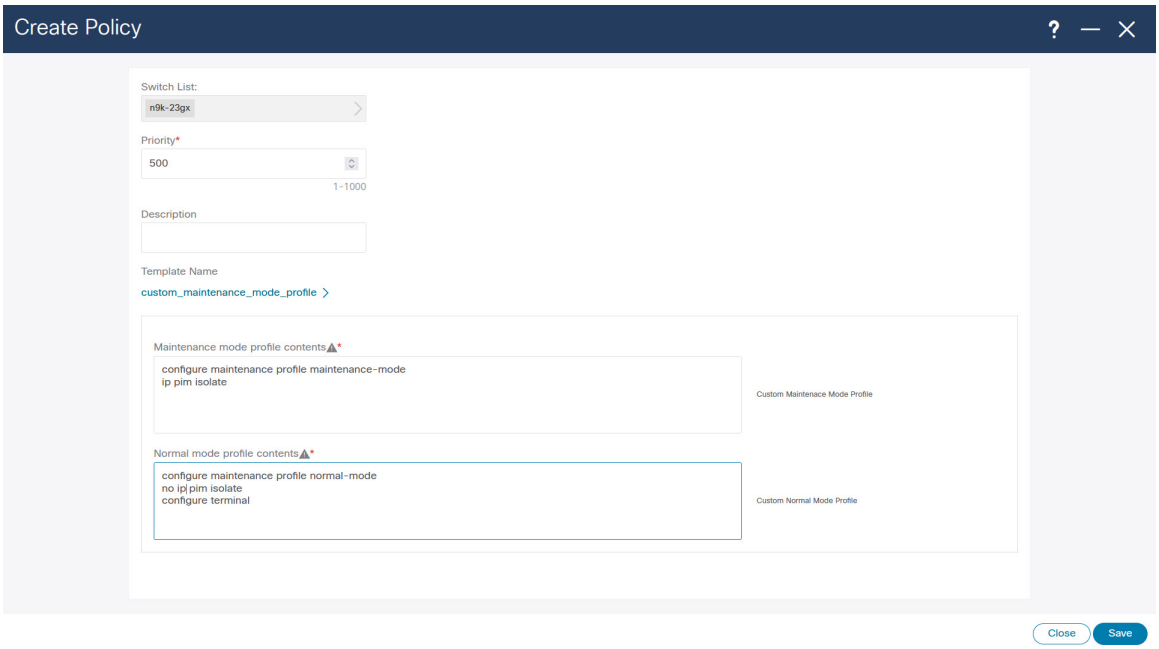
Example:

```
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
```

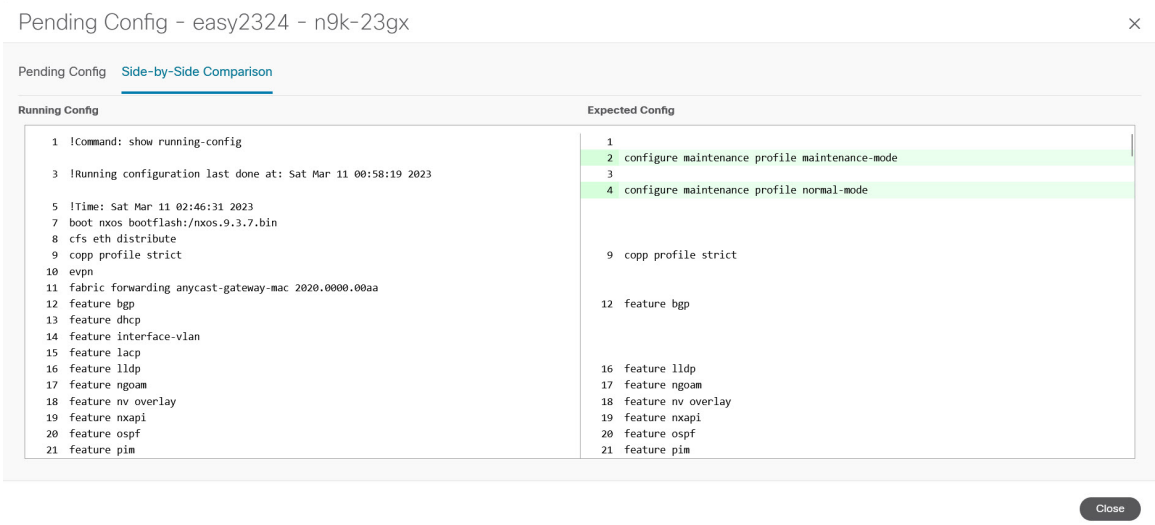



**Step 6** Click **Save**.

**Step 7** From Switch Overview, click **Actions > Preview**.

**Step 8** Click on **Pending Config** lines to view the **Pending Config** and **Side-by-Side Comparison**.

**Step 9** Click **Close**.




**Step 10** From Switch Overview, click **Actions > Deploy**. Click **Deploy All** to deploy the new policy configuration on the switch.

Click **Close** after the deployment is complete.

**Step 11** Select the policy and select **Actions > More > Change Mode**.

**Step 12** In the Mode drop-down list, choose **Maintenance**.

**Step 13** Click **Save and Deploy Now** to move the switch to maintenance mode.

## Deleting Custom Maintenance Mode Profile Policy

The switch has to be moved to active/operational or normal mode before deleting the custom maintenance mode profile policy. To delete a custom maintenance mode profile policy from **Web UI > Switches**, perform the following procedure.

### Procedure

**Step 1** Select the desired switch and launch **Switch Overview**.

**Step 2** From **Switch Overview > Actions > More > Change Mode.**

**Step 3** In the Mode drop-down list, choose **Normal**.

**Step 4** Click **Save and Deploy Now** to move the switch to normal mode.

**Step 5** After the switch has been moved to normal mode, select the **custom_maintenance_mode_profile** policy that has to be deleted.

**Step 6** Choose **Actions > Edit Policy**.

**Step 7** Choose **Actions > Delete Policy** and click **Confirm** to mark the Policy for deletion.

The **Mark Deleted** column shows **true** indicating that the policy is marked for deletion.

**Step 8** Again, choose **Actions > Delete Policy** and click **Confirm** to delete the Policy.

**Step 9** From Switch Overview, click **Actions > Deploy**. Click **Deploy All** to delete the policy configuration on the switch.

Click **Close** after the deployment is complete.

# Event Analytics

Event Analytics includes the following topics:

## Alarms

This tab displays the alarms that are generated for various categories. This tab displays information such as ID (optional), Severity, Failure Source, Name, Category, Acknowledged, Creation Time, Last Updated (optional), Policy, and Message. You can specify the Refresh Interval in this tab. You can select one or more alarms and then acknowledge or unacknowledge their status using the Change Status drop-down list. In addition, you can select one or more alarms and then click the Delete button to delete them.

## Events

This tab displays the events that are generated for the switches. This tab displays information such as Ack, Acknowledged user, Group, Switch, Severity, Facility, Type, Count, Last Seen, and Description. You can select one or more events and then acknowledge or unacknowledge their status using the Change Status drop-down list. In addition, you can select one or more alarms and then click the Delete button to delete them. If you want to delete all events, click the Delete All button.

The following table describes the fields that appear on **Operations > Event Analytics > Events**.

| Field | Description |
|---|---|
| Group | Specifies the Fabric |
| Switch | Specifies the hostname of the switch |
| Severity | Specifies the severity of the event |
| Facility | Specifies the process that creates the events.<br>The event facility includes two categories: NDFC and syslog facility. Nexus Dashboard Fabric Controller facility represents events generated by Nexus Dashboard Fabric Controller internal services and SNMP traps generated by switches. Syslog facility represents the machine process that created the syslog messages. |
| Type | Specifies how the switch/fabric are managed |
| Count | Specifies the number of times the event has occurred |
| Creation Time | Specifies the time when the event was created |
| Last Seen | Specifies the time when the event was run last |
| Description | Specifies the description provided for the event |
| Ack | Specifies if the event is acknowledged or not |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **Operations > Event Analytics > Events**.

| Action Item | Description |
|---|---|
| Acknowledge | Select one or more events from the table and choose **Acknowledge** icon to acknowledge the event information for the fabric.<br>After you acknowledge the event for a fabric, the acknowledge icon is displayed in the Ack column next to the Group. |
| Unacknowledge | Select one or more events from the table and choose **Unacknowledge** icon to acknowledge the event information for the fabric. |
| Delete | Select an event and choose **Delete** to delete the event. |
| Add Suppressor | Select an event and choose **Add Suppressor** to add a rule to the event. You can provide name to the rule. Using the **Scope** options, you can add this rule to all the Fabrics, or particular elements or all elements. |
| Event Setup | Allows you to setup new event. For more information, see Event Setup, on page 379. |

## Accounting

You can view the accounting information on Cisco Nexus Dashboard Fabric Controller Web UI.

The following table describes the fields that appear on **Operations > Event Analytics > Accounting**.

| Field | Description |
|---|---|
| Source | Specifies the source |
| User Name | Specifies the user name. |
| Time | Specifies the time when the event was created |
| Description | Displays the description. |
| Group | Specifies the name of the group. |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **Operations > Event Analytics > Accounting**.

| Action Item | Description |
|---|---|
| Delete | Select a row and choose **Delete** to delete accounting information from the list. |

# VRFs

### UI Navigation

The following options are applicable only for switch fabrics, easy fabrics, and MSD fabrics.

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **VRFs**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **VRFs**.

**Note** Overlay-mode CLI is available only for Easy and eBGP Vxlan Fabrics.

To create overlay VRFs, create VRFs for the fabric and deploy them on the fabric switches. Before attaching or deploying the VRFs, set the overlay mode. For more information on how to choose the overlay mode, refer the Overlay Mode, on page 83 section.

You can view the VRF details in the **VRFs** horizontal tab and VRF attachment details in the **VRF Attachments** horizontal tab.

This section contains the following:

# VRFs

### UI Navigation

The following options are applicable only for switch fabrics, easy fabrics, and MSD fabrics.

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **VRFs** > **VRFs**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **VRFs** > **VRFs**.

Use this tab to create, edit, delete, import, and export VRFs. You can create networks only after creating VRFs except when you use Layer 2 to create networks.

*Table 1: VRF Table Fields and Description*

| Field | Description |
|---|---|
| VRF Name | Specifies the name of the VRF. |
| VRF Status | Specifies whether the status of the VRF deployment as NA, out-of-sync, pending, deployed, and so on. |
| VRF ID | Specifies the ID of the VRF. |

Click the table header to sort the entries in alphabetical order of that parameter.

The following table describes the action items, in the **Actions** drop-down list, that appears on the **VRFs** horizontal tab of the **VRFs** tab in the **Fabric Overview** window.

*Table 2: VRFs Actions and Description*

| Action Item | Description |
|---|---|
| Create | Allows you to create a new VRF. For more information, see Creating VRF, on page 205. |

| Action Item | Description |
|---|---|
| Edit | Allows you to edit the selected VRF.<br>To edit a VRF, select the check box next to the VRF that you want to edit and choose **Edit**. In the **Edit VRF** window, you can edit the parameters and click **Save** to retain the changes or click **Cancel** to discard the changes. |
| Import | Allows you to import VRF information for the fabric.<br>To import VRF information, choose **Import**. Browse the directory and select the `.csv` file that contains the VRF information. Click **Open**. The VRF information is imported and displayed in the **VRFs** tab of the **Fabric Overview** window. |
| Export | Allows you to export VRF information to a `.csv` file. The exported file contains information pertaining to each VRF, including the configuration details that you saved during the creation of VRFs.<br>To export VRF information, choose **Export**. Select a location on your local system directory to store the VRF information from Nexus Dashboard Fabric Controller and click **Save**. The VRF information file is exported to your local directory. The file name is appended with the date and time at which the file was exported.<br>**Note**<br>You can use the exported `.csv` file for reference or use it as a template for creating new VRFs. |
| Delete | Allows you to delete a selected VRF.<br>To delete a VRF, select the check box next to the VRF that you want to delete and choose **Delete**. You can select multiple VRF entries and delete them at the same instance. A warning message appears asking whether you want to delete the VRF(s). Click **Confirm** to delete or click **Cancel** to retain the VRF. A message appears that the selected VRFs are deleted successfully. |

## Creating VRF

### UI Navigation

The following options are applicable only for switch fabrics, Easy fabrics, and MSD fabrics.

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **VRFs** > **VRFs**.
- Choose **LAN** > **Fabrics**. Double-click on the fabric to open **Fabric Overview** > **VRFs** > **VRFs**.

To create VRF from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Procedure

**Step 1** On the **VRFs** tab, click **Actions** > **Create**.

The **Create VRF** window appears.

**Step 2** On **Create VRF**, enter the required details in the mandatory fields. The available fields vary based on the fabric type.

The fields in this window are:

**VRF Name** – Specifies a VRF name automatically or allows you to enter a name. The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).

For MSD Fabrics, the values for VRF or Network is same for the fabric.

**VRF ID** – Specifies the ID for the VRF or allows you to enter an ID for the VRF.

**VLAN ID** – Specifies the corresponding tenant VLAN ID for the network or allows you to enter an ID for the VLAN. If you want to propose a new VLAN for the network, click **Propose VLAN**.

**VRF Template** – A default universal template is auto-populated. This is applicable for leaf switches only.

**VRF Extension Template** – A default universal extension template is auto-populated. 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.

**Step 3** The fields on the **General** tab are:

**VRF VLAN Name** – Enter the VLAN name for the VRF.

**VRF Interface Description** – Enter a description for the VRF interface.

**VRF Description** – Enter a description for the VRF.

**Step 4** Click the **Advanced** tab to optionally specify the advanced profile settings. The fields on this tab are auto-populated. The fields on the **Advanced** tab are:

**VRF Interface MTU** – Specifies VRF interface MTU.

**Loopback 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 also.

**Redistribute Direct Route Map** – Specifies the redistribute direct route map name.

**Max BGP Paths** – Specifies the maximum number of BGP paths. The valid value is between 1 and 64.

**Max iBGP Paths** – Specifies the maximum number of iBGP paths. The valid value is between 1 and 64.

**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.

**TRM Enable** – Check the check box to enable TRM.

If you enable TRM, and provide the RP address, you must enter the underlay multicast address in the **Underlay Mcast Address**.

**Is RP External** – Check this check box if the RP is external to the fabric. If this check box is not checked, 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 TRM BGW MSite** – Check the check box to enable TRM on Border Gateway Multisite.

**Advertise Host Routes** – Check this check box to control advertisement of /32 and /128 routes to Edge routers.

**Advertise Default Route** – Check this check box to control advertisement of default route 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** check box) for the associated VRF. This will result in /32 routes for hosts 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 one fabric only then the default route is sufficient for inter-subnet communication.

**Config Static 0/0 Route** – Check this check box to control configuration of static default route.

**BGP Neighbor Password** – Specifies the VRF Lite BGP neighbor password.

**BGP Password Key Encryption Type** – From the drop-down list, select the encryption type.

**Enable Netflow** – Allows you to enable netflow monitoring on the VRF-Lite sub-interface. Note that this is supported only if netflow is enabled on the fabric.

**Netflow Monitor** – Specifies the monitor for the VRF-lite netflow configuration.

To enable netflow on a VRF-Lite sub-interface, you must enable netflow at VRF level and VRF extension level. Check the **Enable_IFC_Netflow** check box in the VRF attachment while you edit an extension to enable netflow monitoring.

For more information, refer to Netflow Support, on page 149.

**Step 5** Click **Create** to create the VRF or click **Cancel** to discard the VRF.

A message appears indicating that the VRF is created.

The new VRF appears on the **VRFs** horizontal tab. The status is **NA** as the VRF is created but not yet deployed. Now that the VRF is created, you can create and deploy networks on the devices in the fabric.

## VRF Attachments

### UI Navigation

The following options are applicable only for switch fabrics, Easy fabrics, and MSD fabrics.

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **VRFs** > **VRF Attachments**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **VRFs** > **VRF Attachments**.

Use this window to attach or detach attachments to or from a VRF, respectively. You can also import or export the attachments for a VRF.

***Table 3: VRF Attachments Table Fields and Description***

| Field | Description |
|---|---|
| VRF Name | Specifies the name of the VRF. |
| VRF ID | Specifies the ID of the VRF. |
| VLAN ID | Specifies the VLAN ID. |
| Switch | Specifies the name of the switch. |
| Status | Specifies the status of VRF attachments, for example, pending, NA, deployed, out-of-sync, and so on. |
| Attachment | Specifies whether the VRF attachment is attached or detached. |
| Switch Role | Specifies the switch role. For example, for the fabric created using the Easy_Fabric_IOS_XE fabric template, the switch role is specified as either leaf, spine, or border. |
| Fabric Name | Specifies the name of the fabric to which the VRF is attached or detached. |
| Loopback ID | Specifies the loopback ID. |
| Loopback IPV4 Address | Specifies the loopback IPv4 address. |
| Loopback IPV6 Address | Specifies the loopback IPv6 address.<br>**Note**<br>The IPv6 address is not supported for underlay. |

Click the table header to sort the entries in alphabetical order of that parameter.

The following table describes the action items, in the **Actions** drop-down list, that appears on the **VRF Attachments** horizontal tab of the **VRFs** tab in the **Fabric Overview** window.

*Table 4: VRF Attachments Actions and Description*

| Action Item | Description |
|---|---|
| History | Allows you to view the deployment and policy change history of the selected VRF.<br><br>You can view the deployment history details of a VRF attachment such as hostname, VRF name, commands, status, status description, user, and completed time on the **Deployment History** tab.<br><br>You can view the policy change history details such as policy ID, template, description, PTI operation, generated configuration, entity name and type, created date, serial number, user, and source of the policy on the **Policy Change History** tab.<br><br>To view the history of a VRF attachment, check the check box next to the VRF name and select **History**. The **History** window appears. Click the **Deployment History** or **Policy Change History** tabs as required. You can also click the **Detailed History** link in the **Commands** column of the **Deployment History** tab to view the command execution details (comprising configuration, status, and CLI response) for the host. |
| Edit | Allows you to view or edit the VRF attachment parameters such as interfaces that you want to attach to the selected VRF.<br><br>To edit the VRF attachment information, check the check box next to the VRF name that you want to edit. Select **Edit**. In the **Edit VRF Attachment** window, edit the required values, attach or detach the VRF attachment. Click the **Edit** link to edit the CLI freeform config for the switch, and click **Save** to apply the changes or click **Cancel** to discard the changes. The edited VRF attachment is shown in the table on the **VRF Attachments** horizontal tab of the **VRFs** tab in the **Fabric Overview** window. |
| Preview | Allows you to preview the configuration of the VRF attachments for the selected VRF.<br><br>**Note**<br>This action is not allowed for attachments that are in deployed or NA status.<br><br>To preview the VRF, check the check box next to the VRF name and choose **Preview** from **Actions** drop-down list. The **Preview Configuration** window for the fabric appears.<br><br>You can preview the VRF attachment details such as the VRF name, fabric name, switch name, serial number, IP address, and role, VRF status, pending configuration, and progress of the configuration. Click the lines link in the **Pending Config** column to view the lines for which the configuration is pending. Click **Close**. |

| Action Item | Description |
|---|---|
| Deploy | Allows you to deploy the pending configuration of the VRF attachments, for example, interfaces, for the selected VRF.<br>**Note**<br>This action is not allowed for attachments that are in deployed or NA status.<br><br>To deploy a VRF, check the check box next to the VRF name and choose **Deploy** from **Actions** drop-down list. The **Deploy Configuration** window for the fabric appears.<br><br>You can view the details such as the VRF name, fabric name, switch name, serial number, IP address, and role, VRF status, pending configuration, and progress of the configuration. Click the lines link in the **Pending Config** column to view the lines for which the configuration is pending. Click the **Deploy** button. The status and progress of the deployment is displayed in the **VRF Status** and **Progress** columns. After the deployment is completed successfully, close the window. |
| Import | Allows you to import information about VRF attachments for the selected fabric.<br><br>To import the VRF attachments information, choose **Import**. Browse the directory and select the `.csv` file that contains the VRF attachments information. Click **Open** and then click **OK**. The VRF information is imported and displayed in the **VRF Attachments** horizontal tab on the **VRFs** tab in the **Fabric Overview** window. |
| Export | Allows you to export the information about VRF attachments to a `.csv` file. The exported file contains information pertaining to each VRF, including the fabric it belongs to, whether the LAN is attached, the associated VLAN, serial number, interfaces, and freeform configuration details that you saved for VRF attachments.<br><br>To export VRF attachments information, choose the **Export** action. Select a location on your local system directory to store the VRF information and click **Save**. The VRF information file is exported to your local directory. The file name is appended with the date and time at which the file was exported. |
| Quick Attach | Allows you to immediately attach an attachment to the selected VRF. You can select multiple entries and attach them to a VRF at the same instance.<br><br>To quickly attach any attachment to a VRF, choose **Quick Attach** from **Actions** drop-down list. A message appears to inform that the attach action was successful. |

| Action Item | Description |
|---|---|
| Quick Detach | Allows you to detach the selected VRF immediately from an attachment, for example, a fabric. You can select multiple entries and detach them from an attachment at the same instance.<br>To attach any attachment to a VRF quickly, choose **Quick Detach** from **Actions** drop-down list. A message appears to inform that the detach action was successful. |

# Networks

**UI Navigation**

The following options are applicable only for switch fabrics, easy fabrics, and MSD fabrics:

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Networks**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Networks**.

**Note** Before creating networks, ensure that you have created a VRF for the fabric. However, if you have chosen Layer 2, you do not require a VRF. For more information about VRFs, see VRFs, on page 203.

To create overlay networks, create networks for the fabric and deploy them on the fabric switches. Before deploying the networks, set the overlay mode. For more information on how to choose the overlay mode, refer the Overlay Mode, on page 83 section.

For information about creating interface groups and attaching networks, see Interface Groups, on page 330.

You can view the network details in the **Networks** horizontal tab and network attachment details in the **Network Attachments** horizontal tab.

This section contains the following:

## Networks

The following table describes the action items, in the **Actions** drop-down list, that appears in the **Networks** window.

***Table 5: Networks Actions and Description***

| Action Item | Description |
|---|---|
| Create | Allows you to create a new network for the fabric. For instructions about creating a new network, see Creating Network for Standalone Fabrics, on page 214. |

| Action Item | Description |
|---|---|
| Edit | Allows you to view or edit the selected network parameters.<br><br>To edit the network information, select the check box next to the network name that you want to edit and choose **Edit**. In the **Edit Network** window, edit the required values and click **Submit** to apply the changes or click **Cancel** to discard the host alias. The edited network is shown in the table in the **Networks** tab of the **Fabric Overview** window. |
| Import | Allows you to import network information for the fabric.<br><br>To import network information, choose **Import**. Browse the directory and select the `.csv` file that contains the host IP address and corresponding unique network information. Click **Open**. The host aliases are imported and displayed in the **Networks** tab of the **Fabric Overview** window. |
| Export | Allows you to 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.<br><br>To export network information, choose **Export**. Select a location on your local system directory to store the network information from Nexus Dashboard Fabric Controller and click **Save**. The network information file is exported to your local directory. The file name is appended with the date and time at which the file was exported.<br><br>**Note**<br>You can use the exported `.csv` file for reference or use it as a template for creating new networks. Before importing the file, 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. |
| Delete | Allows you to delete the network.<br><br>To delete a network for the fabric, select the check box next to the network name that you want to delete and choose **Delete**. You can select multiple network entries and delete them at the same instance. |

| Action Item | Description |
|---|---|
| Add to interface group | Allows you to add the network to an interface group. You can select multiple network entries and add them to an interface group at the same instance.<br><br>To add the selected networks to the interface group that you want, click **Add to interface group** action.<br><br>In the **Add to interface group** window, click the networks link and verify whether the selected networks are present in the **Selected Networks** window and then close the window. Either select an interface group from the drop-down list or click **Create new interface group**.<br><br>In the **Create new interface group** window, provide the interface group name, select the interface type, and then click **Save** to save the changes and close the window or click **Cancel** to discard the changes.<br><br>In the **Add to interface group** window, click **Save** to save the changes and close the window or click **Cancel** to discard the changes.<br><br>The interface group is displayed in a column in the **Networks** tab of the **Fabric Overview** window. |
| Remove from interface group | Allows you to remove the network from an interface group. You can select multiple network entries and remove them from an interface group at the same instance.<br><br>To remove the selected networks to the interface group that you want, click **Remove from interface group** action.<br><br>In the **Remove from interface group** window, click the networks link and verify whether the selected networks are present in the **Selected Networks** window and then close the window.<br><br>In the **Remove from interface group** window, click **Remove** to remove the networks from the interface group and close the window or click **Cancel** to discard the changes.<br><br>The interface group are removed from the column in the **Networks** tab of the **Fabric Overview** window. |

***Table 6: Networks Table Fields and Description***

| Field | Description |
|---|---|
| Network Name | Specifies the name of the network. |
| Network Id | Specifies the Layer 2 VNI of the network. |
| VRF Name | Specifies the name of the Virtual Routing and Forwarding (VRF). |
| IPv4 Gateway/Suffix | Specifies the IPv4 address with subnet. |

| Field | Description |
|---|---|
| IPv6 Gateway/Suffix | Specifies the IPv6 address with subnet. |
| Network Status | Displays the status of the network. |
| Vlan Id | Specifies the VLAN Id. |
| Interface Group | Specifies the interface group. |

## Creating Network for Standalone Fabrics

To create a network from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Before you begin

Before creating networks, ensure that you have created a VRF for the fabric. However, if you have chosen Layer 2 on the **Create Network** window, then you do not require a VRF. For more information, see VRFs, on page 203.

### Procedure

**Step 1** On the **Networks** tab, click **Actions** > **Create**.

The **Create Network** window appears.

**Step 2** On **Create Network**, enter the required details in the mandatory fields. The available fields vary based on the fabric type.

**Note**
If the fields for the **Network ID** field below and the **VRF ID** field (after clicking **Create VRF**) are not automatically populated, one possible reason is that the VNI ranges might be exhausted. In this situation, you can extend the range for VNI accordingly in **Fabric Settings**.

The fields in this window are:

**Network ID** and **Network Name** – Specifies the Layer 2 VNI and the 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.

**Layer 2 Only** – Specifies whether the network is Layer 2 only.

**VRF Name** – Allows you to select the Virtual Routing and Forwarding (VRF) from the drop-down list.

If you want to create a new VRF, click **Create VRF**. The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).

**VLAN ID** – Specifies the corresponding tenant VLAN ID for the network. If you want to propose a new VLAN for the network, click **Propose VLAN**.

**Network Template** – A default universal template is auto-populated. This is only applicable for leaf switches.

**Network Extension Template** – A default universal extension template is auto-populated. 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.

**Generate Multicast IP** – Click to generate a new multicast group address and override the default value.

**Step 3** The fields on the **General Parameters** tab are:

**Note**
If the network is a non Layer 2 network, then it is mandatory to provide the gateway IP address.

**IPv4 Gateway/NetMask**: Specifies the IPv4 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. The anycast gateway IP address is the same for MyNetwork_30000 on all switches of the fabric that have the presence of the network.

**Note**
If the same IP address is configured in the IPv4 Gateway and IPv4 Secondary GW1 or GW2 fields of the network template, Nexus Dashboard Fabric Controller 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 List** – Specifies the IPv6 address with subnet.

**Vlan Name** – Enter the VLAN name.

**Interface Description** – Specifies the description for the interface. This interface is a switch virtual interface (SVI).

**MTU for L3 interface** – Enter the MTU for Layer 3 interfaces range 68 - 9216.

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

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

**IPv4 Secondary GW3** – Enter the gateway IP address for the additional subnet.

**IPv4 Secondary GW4** – Enter the gateway IP address for the additional subnet.

**Step 4** Click the **Advanced** tab to optionally specify the advanced profile settings. The fields on the **Advanced** tab are:

**ARP Suppression** – Select the check box to enable the ARP Suppression function.

**Ingress Replication** – The check box 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 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.

Starting from Cisco NDFC Release 12.1.2e, a maximum of 16 DHCP relay servers for overlay networks are supported. Perform the following steps to include the DHCP relay server information:

**a.** a. On the **DHCP Relay Server Information** field, click **Actions** > **Add**.

The **ADD Item** window appears.

b. Enter the **Server IP V4 Address** and **Server VRF** details and click **Save**.

c. Repeat the above steps to add the required number of DHCP relay server information.

**DHCPv4 Server 1** – Enter the DHCP relay IP address of the first DHCP server.

**DHCPv4 Server VRF** – Enter the DHCP server VRF ID.

**DHCPv4 Server 2** – Enter the DHCP relay IP address of the next DHCP server.

**DHCPv4 Server2 VRF** – Enter the DHCP server VRF ID.

**DHCPv4 Server 3** – Enter the DHCP relay IP address of the next DHCP server.

**DHCPv4 Server3 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** – Check the check box to enable TRM.

For more information, see Overview of Tenant Routed Multicast.

**L2 VNI Route-Target Both Enable** – Check the check box to enable automatic importing and exporting of route targets for all L2 virtual networks.

**Enable Netflow** – Enables netflow monitoring on the network. This is supported only if netflow is already enabled on fabric.

**Interface Vlan Netflow Monitor** – Specifies the netflow monitor specified for Layer 3 record for the VLAN interface. This is applicable only if **Is Layer 2 Record** is not enabled in the **Netflow Record** for the fabric.

**Vlan Netflow Monitor** – Specifies the monitor name defined in the fabric setting for Layer 3 **Netflow Record**.

**Enable L3 Gateway on Border** – Check the check box to enable a Layer 3 gateway on the border switches.

**Step 5** Click **Create**.

A message appears indicating that the network is created.

The new network appears on the **Networks** page that comes up.

The Status is **NA** since the network is created but not yet deployed on the switches. Now that the network is created, you can create more networks if necessary and deploy the networks on the devices in the fabric.

## Network Attachments

**UI Navigation**

The following options are applicable only for switch fabrics, Easy fabrics, and MSD fabrics:

- Choose **LAN** > **Fabrics**. Click on the fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Networks** > **Network Attachments**.
- Choose **LAN** > **Fabrics**. Double-click on the fabric to open **Fabric Overview** > **Networks** > **Network Attachments**.

Use this window to attach fabrics and interfaces to a network.

*Table 7: Network Attachments Table Fields and Description*

| Field | Description |
|---|---|
| Network Name | Specifies the name of the network. |
| Network ID | Specifies the Layer 2 VNI of the network. |
| VLAN ID | Specifies the VLAN ID. |
| Switch | Specifies the name of the switch. |
| Ports | Specifies the ports for the interfaces. |
| Status | Specifies the status of the network attachments, for example, pending, NA, and so on. |
| Attachment | Specifies whether the network attachment is attached or detached. |
| Switch Role | Specifies the switch role. For example, for the fabric created using the Easy_Fabric_IOS_XE fabric template, the switch role is specified as either leaf, spine, or border. |
| Fabric Name | Specifies the name of the fabric to which the network is attached or detached. |

The following table describes the action items, in the **Actions** drop-down list, that appears in the **Network Attachments** horizontal tab on the **Networks** tab in the **Fabric Overview** window.

*Table 8: Network Attachments Actions and Description*

| Action Item | Description |
|---|---|
| History | Allows you to view the deployment and policy change history of the selected network.<br><br>You can view the deployment history details of a network attachment such as hostname, network name, VRF name, commands, status, status description, user and completed time on the **Deployment History** tab.<br><br>You can view the policy change history details such as policy ID, template, description, PTI operation, generated configuration, entity name and type, created date, serial number, user, and source of the policy on the **Policy Change History** tab.<br><br>To view the history of a network attachment, select the check box next to the network name and choose the **History** action. The **History** window appears. Click the **Deployment History** or **Policy Change History** tabs as required. Click the **Detailed History** link in the **Commands** column of the **Deployment History** tab to view the command execution details (comprising configuration, status, and CLI response) for the host. |

| Action Item | Description |
|---|---|
| Edit | Allows you to view or edit the network attachment parameters such as interfaces that you want to attach to the selected network.<br><br>To edit the network attachment information, check the check box next to the network name that you want to edit and choose the **Edit** action. In the **Edit Network Attachment** window, edit the required values, attach or detach the network attachment, click the **Edit** link to edit the CLI freeform config for the switch, and click **Save** to apply the changes or click **Cancel** to discard the changes. The edited network attachment is shown in the table on the **Network Attachments** horizontal tab of the **Networks** tab in the **Fabric Overview** window. |
| Preview | Allows you to preview the configuration of the network attachments for the selected network.<br><br>**Note**<br>This action is not allowed for attachments that are in deployed or NA status.<br><br>To preview the network, check the check box next to the network name and choose **Preview** from **Actions** drop-down list. The **Preview Configuration** window for the fabric appears.<br><br>You can preview the network attachment details such as the network name, fabric name, switch name, serial number, IP address, and role, network status, pending configuration, and progress of the configuration. Click the lines link in the **Pending Config** column to view the lines for which the configuration is pending. Click **Close**. |
| Deploy | Allows you to deploy the pending configuration of the network attachments, for example, interfaces, for the selected network.<br><br>**Note**<br>This action is not allowed for attachments that are in deployed or NA status.<br><br>To deploy a network, check the check box next to the network name and choose **Deploy** from **Actions** drop-down list. The **Deploy Configuration** window for the fabric appears.<br><br>You can view the details such as the network name, fabric name, switch name, serial number, IP address, and role, network status, pending configuration, and progress of the configuration. Click the lines link in the **Pending Config** column to view the lines for which the configuration is pending. Click the **Deploy** button. The status and progress of the deployment is displayed in the **Network Status** and **Progress** columns. After the deployment is completed successfully, close the window. |

| Action Item | Description |
|---|---|
| Import | Allows you to import information about network attachments for the selected fabric.<br><br>To import the network attachments information, choose **Import**. Browse the directory and select the `.csv` file that contains the network attachments information. Click **Open** and then click **OK**. The network information is imported and displayed in the **Network Attachments** horizontal tab on the **Networks** tab in the **Fabric Overview** window. |
| Export | Allows you to export the information about network attachments to a `.csv` file. The exported file contains information pertaining to each network, including the fabric it belongs to, whether the LAN is attached, the associated VLAN, serial number, interfaces, and freeform configuration details that you saved for network attachments.<br><br>To export network attachments information, choose the **Export** action. Select a location on your local system directory to store the network information and click **Save**. The network information file is exported to your local directory. The file name is appended with the date and time at which the file was exported. |
| Quick Attach | Allows you to immediately attach an attachment to the selected network. You can select multiple entries and attach them to a network at the same instance.<br><br>**Note**<br>Interfaces cannot be attached to a network using this action.<br><br>To quickly attach any attachment to a network, choose **Quick Attach** from **Actions** drop-down list. A message appears to inform that the attach action was successful. |
| Quick Detach | Allows you to immediately detach the selected network from an attachment, for example, a fabric. You can select multiple entries and detach them from an attachment at the same instance.<br><br>To quickly detach any attachment to a network, choose **Quick Detach** from **Actions** drop-down list. A message appears to inform that the detach action was successful.<br><br>After quick detach, the switch status is not computed when there is no deploy. Post deploy, the configuration compliance calls at entity level (interface or overlay). |

# History

The history tab displays information about the deployment and policy change history. Choose **LAN > Fabrics**. Double-click a fabric name to open the **Fabric Overview** window and then click the **History** tab.

## Viewing Deployment History

Deployment history of the switches and networks that are involved in the selected service policy or route peering are displayed in the **Deployment History** tab. The deployment history captures the changes that are pushed or deployed from Nexus Dashboard Fabric Controller to switches. The deployment history captures the changes that are pushed or deployed from Nexus Dashboard Fabric Controller to switches.

The following table describes the fields that appear on this page.

| Field | Description |
|---|---|
| Hostname(Serial Number) | Specifies the host name. |
| Entity Name | Specifies the entity name. |
| Entity Type | Specifies the entity type. |
| Source | Specifies the source. |
| Commands | Specifies the commands. |
| Status | Specifies the status of the host. |
| Status Description | Specifies the status description. |
| User | Specifies the user. |
| Time of Completion | Specifies the timestamp of the deployment. |

## Viewing Policy Change History

Different users can simultaneously change expected configuration of switches in the Nexus Dashboard Fabric Controller. You can view the history of policy changes in the **Policy Change History** tab.

The following table describes the fields that appear on this page.

| Field | Description |
|---|---|
| Policy ID | Specifies the policy ID. |
| Template | Specifies the template that is used. |
| Description | Specifies the description. |
| PTI Operation | Specifies the Policy Template Instances (PTIs). |
| Generated Config | Specifies the configuration history. Click **Detailed History** to view the configuration history. |
| Entity Name | Specifies the entity name. |
| Entity Type | Specifies the entity type. |
| Created On | Specifies that date on which the policy was created. |
| Priority | Specifies the priority value. |

| Field | Description |
|---|---|
| Serial Number | Specifies the serial number. |
| Content Type | Specifies the content type. |
| User | Specifies the user. |
| Source | Specifies the source. |

# Resources

Cisco Nexus Dashboard Fabric Controller allows you to manage the resources. The following table describes the fields that appear on this page.

| Field | Description |
|---|---|
| Scope Type | Specifies the scope level at which the resources are managed. The scope types can be **Fabric**, **Device**, **Device Interface**, **Device Pair**, and **Link**. |
| Scope | Specifies the resource usage scope. Valid values are the switch serial numbers or fabric names. Resources with serial numbers are unique and can be used on the serial number of the switch only. |
| Device Name | Specifies the name of the device. |
| Device IP | Specifies the IP address of the device. |
| Allocated Resource | Specifies if the resources are managed with device, device interface, or fabric. Valid values are ID type, subnet, or IP addresses. |
| Allocated To | Specifies the entity name for which the resource is allocated. |
| Resource Type | Specifies the resource type. The valid values are **TOP_DOWN_VRF_LAN**, **TOP_DOWN_NETWORK_VLAN**, **LOOPBACK_ID**, **VPC_ID**, and so on. |
| Is Allocated? | Specifies if the resource is allocated or not. The value is set to **True** if the resource is permanently allocated to the given entity. The value is set to **False** if the resource is reserved for an entity and not permanently allocated. |
| Allocated On | Specifies the date and time of the resource allocation. |
| ID | Specifies the ID. |

## Allocating a Resource

To allocate a resource from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Procedure

**Step 1** Choose **LAN > Fabrics**.

**Step 2** Double-click a fabric name.

The **Fabric Overview** window appears.

**Step 3** Click the **Resources** tab.

**Step 4** Click **Actions > Allocate Resource** to allocate the resource.

The **Allocate Resource** window appears.

**Step 5** Choose the pool type, pool name, and scope type from the drop-down lists accordingly.

The options for pool type are **ID_POOL**, **SUBNET_POOL**, and **IP_POOL**. Based on the pool type you choose, the values in the **Pool Name** drop-down list changes.

**Step 6** Enter the entity name in the **Entity Name** field.

The embedded help gives example names for different scope types.

**Step 7** Enter the ID, IP address, or the subnet in the **Resource** field based on what pool type you chose in *Step 3*.

**Step 8** 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 Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Procedure

**Step 1** Choose **LAN > Fabrics**.

**Step 2** Double-click a fabric name.

The **Fabric Overview** window appears.

**Step 3** Click the **Resources** tab.

**Step 4** Choose a resource that you want to delete.

**Note**
You can delete multiple resources at the same time by choosing multiple resources.

**Step 5** Click **Actions > Release Resource(s)** to release the resource.

A confirmation dialog box appears.

**Step 6** Click **Confirm** to release the resource.

# Hosts



**Note** This tab is only available on IPFM fabric when you have deployed IPFM on Nexus Dashboard Fabric Controller.

**Nexus Dashboard Fabric Controller UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Hosts**.

Information about hosts is also displayed as a card on the **Overview** tab in the **Fabric Overview** window. For more information about these cards, see Hosts, on page 183.

The **Hosts** tab includes the following tabs:

## Discovered Hosts Summary

**Nexus Dashboard Fabric Controller UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Hosts** > **Discovered Hosts Summary**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Hosts** > **Discovered Hosts Summary**.

You can view a summary of all the hosts that are populated through telemetry in this window.

***Table 9: Discovered Hosts Summary Table Fields and Description***

| Field | Description |
|---|---|
| VRF | Specifies the VRF for the host. |
| Host | Specifies the IP address for the host. |
| Senders/Receivers | Specifies the number of times the host device plays its role as a sender or a receiver. Click the count to view where it was used. |

Click the table header to sort the entries in alphabetical order of that parameter.

## Discovered Hosts

**Nexus Dashboard Fabric Controller UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Hosts** > **Discovered Hosts**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Hosts** > **Discovered Hosts**.

You can view all the hosts that are populated through telemetry on this screen. After the switches are discovered, all the switches in the fabric will push data to the Nexus Dashboard Fabric Controller server at regular intervals using telemetry. Cisco Nexus Dashboard Fabric Controller server displays the received Events and Flow statistics for each active flow.

***Table 10: Discovered Hosts Table Fields and Description***

| Field | Description |
|---|---|
| VRF | Specifies the VRF for the host. |
| Host | Specifies the IP address for the host. |
| Role | Specifies the role of the host device. The role of the host can be one of the following:<br>• Sender<br>• External Sender<br>• Dynamic Receiver<br>• External Receiver<br>• Static Receiver |
| Multicast Group | Specifies the multicast address of the flow in which the host participates. |
| Source | Specifies the source of the flow which the discovered host participates in. |
| Switch | Specifies the name of the switch. |
| Interface | Specifies the interface to which the host is connected to on the sender or receiver switch. |
| MAC Address | Specifies the MAC address of a physical host, if the switch has ARP entry for that host). |
| Host Discovered Time | Specifies the date and time at which the switch discovered the host. |
| Fault Reason | Specifies the failure reason for the flow that the discovered host has participates in. |

Click the table header to sort the entries in alphabetical order of that parameter.

## Host Policies

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric name to open the **Fabric** slide-in pane. Click the Launch icon. Choose **Fabric Overview** > **Hosts** > **Host Policies**.

- Choose **LAN** > **Fabrics**. Double-click on a fabric name to open **Fabric Overview** > **Hosts** > **Host Policies**.

You can add policies to the host devices. Navigate to **Host Policies** to configure the host policies.

**Note** Switches must be deployed with default host policies. You can edit the default host policies to permit or deny. From the Deployment drop-down list, select **Deploy Selected Policies** to deploy the default policies to the switches. You can also deploy all the default policies to all the managed switches by selecting **Deploy All Default Policies** even without selecting any default policies.

By default, the sequence numbers for policies are auto-generated by Nexus Dashboard Fabric Controller and Multicast mask/prefix is taken as /32. If you want to enter the required values for the sequence number and the multicast mask/prefix in the appropriate fields, ensure that the **Enable mask/prefix for the multicast range in Host Policy** check box under **Settings > Server Settings > IPFM** tab is enabled. Then, you can enter the sequence number and the multicast mask/prefix in the appropriate fields available in the **Create Host Policy** and **Edit Host Policy** options available in the **Actions** drop-down list in the **Host Policies** window.

The default host policies must be deployed successfully to the switch before you deploy the custom host policies on that switch. Otherwise, the custom policies will fail to deploy. Ensure that you deploy all default policies successfully to all the switches before you create, edit, import, or deploy custom policies.

**Note** When a user logs in to Nexus Dashboard Fabric Controller with a network operator role, all the buttons or options to create, delete, edit, import, export, or deploy policies are disabled. This user can only view policies, and deployment status or history.

Policies are automatically deployed to switches whenever they are created, edited, or imported. You can choose to undeploy or redeploy the policies, by selecting one or more check boxes next to the policies and choosing the appropriate actions in the **Actions** drop-down list. The policies will not be deployed correctly if the device is rebooting while the policy was deployed. In such case, the Failed message appears in the **Deployment Status** column in the **Host Policies** window.

**Note** If you have created a custom or non-default VRF, although the host and flow policies are automatically created for the VRF, use the action options in this window to manually deploy the host policies to the switches in the fabric.

The following table describes the action items, in the **Actions** drop-down list, that appears in the **Host Policies** window.

***Table 11: Host Policies Actions and Description***

| Action Item | Description |
|---|---|
| Create Host Policy | Allows you to create a new host policy. For instructions about creating a host policy, see Create Host Policy, on page 231. |

| Action Item | Description |
| --- | --- |
| Edit Host Policy | Allows you to view or edit the selected host policy parameters.<br><br>To edit the host policy, select the check box next to the host policy that you want to delete and choose **Edit Host Policy**. In the **Edit Host Policy** window, edit the required values and click **Save & Deploy** to configure and deploy the policy or click **Cancel** to discard the host policy. The edited host policy is shown in the table in the **Host Policies** window.<br><br>**Note**<br>The changes made to host policy are applied immediately. If the policy is already applied to any device, the changes may impact the existing flows. |
| Delete Host Policy | Allows you to delete user-defined host policies.<br><br>**Note**<br>• Undeploy policies from all the switches before deleting them from Nexus Dashboard Fabric Controller.<br>• Default policy can be undeployed from the switches on which it is deployed. However, Custom policy can be deleted and undeployed.<br>• When you undeploy the default policies, all default policies are reset to have default permission (Allow).<br><br>To delete a host policy, select the check box next to the host policy that you want to delete and choose **Delete Host Policy**. You can select multiple host policy entries and delete them at the same instance.<br><br>A delete host policy successful message appears at the bottom of the page. |
| Purge | Allows you to delete all custom policies without selecting any policy check box.<br><br>**Note**<br>• Undeploy policies from all switches before deleting them from Nexus Dashboard Fabric Controller.<br>• You can undeploy the default policy, but you cannot delete the default policy. You can delete and undeploy only the custom policies. |

| Action Item | Description |
|---|---|
| Import | Allows you to import host policies from a CSV file to Nexus Dashboard Fabric Controller.<br>**Note**<br>After import, all policies imported from a CSV file are applied to all managed switches automatically.<br>To import a host policies, choose **Import**. Browse the directory and select the `.csv` file that contains the host policy configuration information. The policy will not be imported if the format in the `.csv` file is incorrect. Click **Open**. The imported policies are automatically deployed to all the switches in the fabric. |
| Export | Allows you to export host policies from Nexus Dashboard Fabric Controller to a `.csv` file.<br>To export host policies, choose **Export**. Select a location on your local system directory to store the host policy details file. Click **Save**. The host policy file is exported to your local directory. The filename is appended with the date on which the file is exported. The format of the exported file is `.csv`. |
| Deploy Selected Policies | Select this option to deploy only the selected policies to the switch. |
| Deploy All Custom Policies | Select this option to deploy all the custom or user-defined policies to the switch in a single instance. If the policies are deployed when the switch is rebooting, the deployment fails and a failed status message appears. |
| Deploy All Default Policies | Select this option to deploy all default policies to the switch. |
| Undeploy Selected Policies | Select this option to undeploy the selected policies.<br>Select one or more check boxes next to the policy name. Select this option from the drop-down list to undeploy the selected policies. |
| Undeploy All Custom Policies | Select this option to undeploy all the custom or user-defined policies in a single instance. |
| Undeploy All Default Policies | Select this option to undeploy the default policies. |
| Redo All Failed Policies | The deployment of policies may fail due to various reasons. Select this option to deploy or undeploy all failed policies.<br>All the deployments that failed on switches previously will be deployed again to only those switches. All the undeployments that failed on switches previously will be undeployed again from only those switches. |

| Action Item | Description |
|---|---|
| Deployment History | Select one policy from the drop-down list. Select this option to view the deployment history of the selected policy in the **Deployment History** pane.<br>The policy name is shown in the Policy Name field. From the drop-down list, choose the switch on which this policy was deployed.<br>The **Deployment History** pane displays the following fields.<br>• Policy Name - Specifies the selected policy name.<br>• VRF - Specifies the VRF for the selected policy.<br>• Switch Name - Specifies the name of the switch that the policy was deployed to.<br>• Deployment Status - Displays the status of deployment. It shows if the deployment was a success, failed, or not deployed. Click on the deployment status, for example, **Success**, to see more details. For more information about the deployment status, see Deployment Status, on page 230.<br>• Action - Specifies the action that is performed on the switch for that host policy. **Create** implies that the policy has been deployed on the switch. **Delete** implies that the policy has been undeployed from the switch.<br>• Deployment Date/Time - Specifies the date and time at which the host policy was last updated. The format is *Day MMM DD YYYY HH:MM:SS Timezone*.<br>• Failed Reason - Specifies why the policy was not successfully deployed. |

***Table 12: Host Policies Table Field and Description***

| Field | Description |
|---|---|
| VRF | Specifies the VRF for the host. The fields—Deployment, Undeployment, Status, and History—are based on VRF. |
| Policy Name | Specifies the policy name for the host, as defined by the user. |
| Receiver | Specifies the IP address of the receiving device. |
| Multicast IP/Mask | Specifies the multicast IP address for the host. |
| Sender | Specifies the IP Address of the transmitting device. |

| Field | Description |
|---|---|
| Host Role | Specifies the host device role. The host device role is either one of the following:<br>• **Sender**<br>• **Receiver**<br>• **Receiver-External**<br>• **Receiver-Local** |
| Operation | Specifies if the operation of the host policy. The policy has the following operations:<br>• **Permit**<br>• **Deny** |
| Sequence Number | Specifies the sequence number of the custom policy when the multicast range is selected. |
| Deployment Action | Specifies the action performed on the switch for that host policy.<br>• **Create** - The policy is deployed on the switch.<br>• **Delete** - The policy is undeployed from the switch. |
| Deployment Status | Specifies if the deployment is successful, failed, or the policy is not deployed. |
| Last Updated | Specifies the date and time at which the host policy was last updated.<br>The format is *Day MMM DD YYYY HH:MM:SS Timezone*. |

### Deployment Status

The following table describes the fields that appear on the Deployment Status.

***Table 13: Deployment Status Field and Description***

| Field | Description |
|---|---|
| Policy Name | Specifies the name of the host policy. |
| VRF | Specifies the name of the VRF. |
| Switch Name | Specifies the switch on which the VRF is deployed. |
| IP Address | Specifies the IP address of the switch. |
| Deployment Status | Displays the status of the deployment. It shows if the deployment was a **Success** or **Failed** along with the reason why the deployment failed. |

| Field | Description |
|---|---|
| Action | Specifies the action that is performed on the switch, for example, **Create**. |
| Deployment Date/Time | Displays the date and time when the deployment was initialized. |

This section contains the following:

## Create Host Policy

### UI Navigation

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Hosts** > **Host Policies**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Hosts** > **Host Policies**.

The default host policies must be deployed successfully to the switch before you deploy the custom host policies on that switch. Otherwise, the custom policies will fail to deploy. Ensure that you deploy all default policies successfully to all the switches before you add custom policies.

To create a host policy from the Cisco Nexus Dashboard Fabric Controller, perform the following steps:

### Procedure

**Step 1** In the **Host Policies** window, from the **Actions** drop-down list, choose **Create Host Policy**.

**Step 2** In the **Create Host Policy** window, specify the parameters in the following fields.

- **VRF** - Click the **Select a VRF** link to open the **Select a VRF** window. The default VRF is also listed in the window. Search and select a VRF for the host and click **Save**.

  **Note**
  - Policy names can be repeated across VRFs, that is, they are unique only within a VRF.
  - Across the VRF, host policies may be same or different.

- **Policy Name** - Specifies a unique policy name for the host policy.
- **Host Role** - Specifies the host as a multicast sender or receiver. Select one of the following:
  - Sender
  - Receiver-Local
  - Receiver-External
- **Sender Host Name** - Specifies the sender host to which the policy is applied.

  **Note**
  Hosts that are discovered as remote senders can be used for creating sender host policies.

- **Sender IP** - Specifies the IP address of the Sender host. Note that you can specify wildcard for this IP address by specifying the ***** (asterisk) symbol or **0.0.0.0** in this field.
- **Receiver Host Name** - Specifies the receiver host to which the policy is applied. If a destination host is detected, you can choose the hostname from the drop-down list.

  **Note**
  Do not select hosts that are discovered as remote receivers to create receiver or sender host policies. However, hosts that are discovered as remote senders can be used for creating sender host policies.
- **Receiver IP** - Specifies the IP address of the receiver host. This field is visible and is applicable only if the Host Role is set to **Receiver-Local**. Note that you can specify wildcard for this IP address by specifying the ***** (asterisk) symbol or **0.0.0.0** in this field.

  **Note**
  When **Receiver IP** in a receiver host policy is a wildcard (***** or **0.0.0.0**), **Sender IP** also has to be a wildcard (***** or **0.0.0.0**).
- **Multicast** - Specifies the multicast IP Address for the host policy. Note that you can specify wildcard for this IP address by specifying the ***** (asterisk) symbol in this field. This will translate to **224.0.0.0/4**. If you specify a wildcard IP address for **Sender IP** and **Receiver IP** fields, the Multicast Group is always required, that is, you cannot specify multicast as ***** or **0.0.0.0**.
- **Permit/Deny** - Click **Permit** if the policy must allow the traffic flow. Click **Deny** if the policy must not allow the traffic flow.

**Step 3** Click **Save & Deploy** to configure and deploy the Policy. Click **Cancel** to discard the new policy. The deployment completed message appears at the bottom of the window. You can click **Refresh** to refresh the current deployment status in the window or click **View Details** to verify the deployment details.

## Host Alias

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Hosts** > **Host Alias**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Hosts** > **Host Alias**.

**Note** This section is applicable for both the IPFM and Generic Multicast modes in Nexus Dashboard Fabric Controller.

Cisco Nexus Dashboard Fabric Controller allows you to create host aliases for sender and receiver hosts for IPFM fabrics. The active multicast traffic transmitting and receiving devices are termed as hosts. You can add a host-alias name to your sender and receiver hosts, to help you identify the hosts by a name. You can also import many Host Alias to Cisco Nexus Dashboard Fabric Controller with IPFM deployment.

The following table describes the action items, in the **Actions** drop-down list, that appears in the **Host Alias** window.

*Table 14: Host Alias Actions and Description*

| Action Item | Description |
|---|---|
| Create Host Alias | Allows you to create a new host alias. For instructions about creating a new host alias, see Create Host Alias, on page 233. |
| Edit Host Alias | Allows you to view or edit the selected host alias parameters.<br>To edit the host alias, select the check box next to the host alias that you want to delete and choose **Edit Host Alias**. In the **Edit Host Alias** window, edit the required values and click **Submit** to apply the changes or click **Cancel** to discard the host alias. The edited host alias is shown in the table in the **Host Alias** window. |
| Delete Host Alias | Allows you to delete the host alias.<br>To delete a host alias, select the check box next to the host alias that you want to delete and choose **Delete Host Alias**. You can select multiple host alias entries and delete them at the same instance. |
| Import | Allows you to import host aliases for devices in the fabric.<br>To import host aliases, choose **Import**. Browse the directory and select the `.csv` file that contains the host IP address and corresponding unique host name information. Click **Open**. The host aliases are imported and displayed in the **Host Alias** window. |
| Export | Allows you to export host aliases for devices in the fabric.<br>To export a host alias, choose **Export**. Select a location on your local system directory to store the host aliases configuration from Nexus Dashboard Fabric Controller and click **Save**. The host alias configuration file is exported to your local directory. The file name is appended with the date and time at which the file was exported. The format of the exported file is `.csv`. |

*Table 15: Host Alias Table Field and Description*

| Field | Description |
|---|---|
| VRF | Specifies the VRF for the host. |
| Host Alias | Specifies the host name that is configured to identify the host. |
| IP Address | Specifies the IP address of the host connecting to the switch, which you want to refer with an alias name. |
| Last Updated At | Specifies the date and time at which the host alias was last updated. |

This section contains the following:

## Create Host Alias

### UI Navigation

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Hosts** > **Host Alias**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Hosts** > **Host Alias**.

Perform the following task to create new host aliases to devices in the fabric discovered by Cisco Nexus Dashboard Fabric Controller.

To create a host alias from the Cisco Nexus Dashboard Fabric Controller, perform the following steps:

**Procedure**

**Step 1** In the **Host Alias** window, from the **Actions** drop-down list, choose **Create Host Alias**.

**Step 2** In the **Create Host Alias** window, enter the following:

**Note**
All the fields are mandatory.

- **VRF** - Select the VRF from this drop-down list. The default value is **default**.

  **Note**
  Host and IP Address are unique per VRF, that is, same host name with the same IP Address can exist in multiple VRFs.

- **Host Name** - Enter a fully qualified unified hostname for identification.
- **IP Address** - Enter the IP address of the host that is part of a flow.

  **Note**
  You can also create host alias before a host sends any data to its directly connected sender or receiver leaf.

**Step 3** Click **Submit** to apply the changes.

Click **Cancel** to discard the host alias.

The new host alias is shown in the table in the **Host Alias** window.

## Applied Host Polices

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Hosts** > **Applied Host Policies**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Hosts** > **Applied Host Policies**.

You can view the policies that you have applied in the entire network on this tab.

The table displays default PIM policy, local receiver policy, and sender policy. IPFM does not display user-defined PIM Policies or Receiver External Policies.

*Table 16: Applied Host Policies Table Fields and Description*

| Column Name | Description |
|---|---|
| VRF | Specifies the VRF for the host. |
| Policy Name/Sequence # | Specifies the name of the policy applied. |
| Host Role | Specifies the role of the host.<br>The host device role is either one of the following:<br>• **PIM**<br>• **Sender**<br>• **Receiver** |
| Switch | Specifies the name of the switch to which the policy is applied. |
| Interface | Specifies the interface to which the policy is applied. |
| Active | Specifies if the policy is active or not. |
| Time Stamp | Specifies the date and time at which the policy was created\deployed.<br>The format is Day, MMM DD YYYY HH:MM:SS (Timezone). |

# Flows

**Note** This tab is only available on IPFM fabric when you have deployed IPFM on Nexus Dashboard Fabric Controller.

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric Summary** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Hosts**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Hosts**.

Information about flows is also displayed as a card on the **Overview** tab in the **Fabric Overview** window. For more information about these cards, see Flows, on page 183.

The **Flows** tab comprises the following horizontal tabs:

## Flow Status

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Hosts** > **Flow Status**.

- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Hosts** > **Flow Status**.

**Note** This section is applicable for both the IPFM and Generic Multicast modes in Nexus Dashboard Fabric Controller.

Cisco Nexus Dashboard Fabric Controller allows you to view the flow status pictorially and statistically.

In the generic multicast mode, switch reports the receiver interface IP address instead of the receiver endpoint IP address. This IP is displayed in the **Flow Status** and **Topology** windows as a host. In the Sender and Receiver fields, the IPs are suffixed with a blue dot and the word **Remote** to indicate that those IPs are remote hosts. Also, as there's no policing of the traffic, switch reports only "allowed bytes/packets" and not "denied bytes/packets".

**Note** To view details for a given flow such as all pre/post multicast and source IP-Addresses, post group, post S/DST ports, pre/post NAT policy ID, starting and destination node details, as well as view the topology, click the **active** hyperlink in **Flow Link State** for a particular multicast IP.

From Cisco NDFC Release 12.1.2e, VXLAN EVPN Tenant Routed Multicast (TRM) flows can be visualized in the **Flows** tab. The flow visualization is for only monitoring existing TRM flows. Clicking on the **active** link displays the end-to-end flow topology.

In VXLAN TRM, Sources and Receivers associated with an overlay flow are in a customer aka tenant VRF. This tenant traffic is encapsulated in an underlay header which has **Encap Source** and **Encap group** (located in default VRF) on the sender VTEP side. The underlay encapsulated flow then reaches the receiver VTEP and is decapsulated here.

The flow topology in NDFC shows the overlay and underlay parts of the flow in different color coding. (purple for underlay and green for overlay).

The following table provides information about the separation between default and tenant VRF and their descriptions:

***Table 17: Active Flow Field and Description***

| Field | Description |
|---|---|
| Type | Specifies the name of the VRF. |
| L3VNI | Specifies the tenant VNI. |
| Encap Source | Specifies the IP address of the encap source from the default VRF. |
| Encap Group | Specifies the IP address of the encap group from the default VRF. |

### Multicast NAT Visualization

Nexus Dashboard Fabric Controller follows the existing flow classification for multicast flows, that is, active, inactive, sender only, or receiver only. With ingress and egress NAT multiple, input and output addresses can

be translated to same group. Nexus Dashboard Fabric Controller aggregates these flows per sender and receiver combination and provides visibility into NAT rules through topology. For more information about flow topology for active flows, see RTP/EDI Flow Monitor, on page 265.

Multicast NAT is supported in the IPFM network, and it is not supported for regular or generic multicast.

You can use the **NAT Search** field to search for NAT flows. All pre/post multicast and source IP-Addresses are not visible in the **Flow Status** window. You can view these details for a given flow in a pop-up by clicking the active flow hyperlink. The **NAT Search** feature allows you to enter the IP address of either pre or post source/multicast group and filter relevant entries. Note that searched IP address may not be visible in main table on filtering as it may be part of pre or post entry that can be seen on corresponding pop-up window.

For NAT flow with NAT type containing Ingress, the source and group will be the post NAT source and post NAT group. For NAT type containing Egress, the source and group will be pre-NAT source and pre-NAT group. NAT rules are displayed on the **Sender Only** and **Receiver Only** tabs.

For a NAT flow, the topology graph path tracing shows the **NAT** badge on the switch which has ingress NAT and shows **NAT** label on the link to the receiver for egress NAT.

For NAT flow, there is an extra table shown below the topology graph panel to show all the relevant Ingress NAT or Egress NAT information. The NAT Flow information is also available on the **Topology** window. This information is available when you click the links in the **Flow Link State** column.

The VRF name is also shown in the slide-in pane for the host and the switch.

For example, **sanjose-vrf:2.2.2.2** indicates that the VRF is sanjose-vrf and the host is 2.2.2.2.

The flows carry the VRF name as prefix. If the VRF is **default**, it will not be displayed.

The following table provides information about the NAT fields and their descriptions:

***Table 18: NAT Field and Description***

| Field | Description |
|---|---|
| NAT | Specifies the NAT mode, that is, Ingress, Egress, or Ingress and Egress.<br>For the Ingress NAT type, the following information is displayed:<br>Ingress (S) – Specifies that ingress NAT is performed on the Sender Switch, also known as First Hop Router (FHR).<br>Ingress (R) - Specifies that ingress NAT is performed on the Receiver Switch (also known as Last Hop Router (LHR).<br>Ingress (S, R) - Specifies that ingress NAT is performed on both the Sender and Receiver Switch. |
| Pre-Source | Specifies the source IP address before NAT. |
| Post-Source | Specifies the source IP address after NAT. |
| Pre-Group | Specifies the multicast group before NAT. |
| Post-Group | Specifies the multicast group after NAT. |
| Post S Port | Specifies the source port after NAT. |
| Post DST Port | Specifies the destination port after NAT. |

The following table describes the fields that appear on the **Active** tab.

*Table 19: Active Tab Fields and Descriptions*

| Field | Description |
|---|---|
| **Common Fields for IPFM and Generic Multicast Modes** | |
| VRF | Specifies the name of the VRF for the flow. |
| Encap | Specifies the name of the encap for the TRM flow. |
| Multicast IP | Specifies the multicast IP address for the flow.<br>**Note**<br>You can click the wave link next to the Multicast IP address to view the pictorial representation of flow statistics. |
| Flow Alias | Specifies the name of the Flow Alias. |
| Flow Link State | Specifies the state of the flow link.<br>Click the **active** link to view the network diagram or topology of the Sender and Receiver.<br>The dotted line displays the direction of the flow of traffic. You can hover over the nodes to view information. The table on the right side shows information about the Sender and Receiver.<br>The flows in the network diagram or topology show the multicast IP as well as the VRF. If the VRF is **default**, then the VRF will not be shown along with the multicast IP. |
| Sender | Specifies the IP Address or the Host alias of the sender for the multicast group. |
| NAT | Specifies whether the flow is ingress, egress, or both ingress and egress. |
| Sender Switch | Specifies if the Sender switch is a leaf or spine. |
| Sender Interface | Specifies the interface to which the sender is connected to. |
| Receiver Switch | Specifies if the Receiver switch is a leaf or spine. |
| Receiver Interface | Specifies the interface to which the receiver is connected to. |
| Sender Start Time | Displays the time from when the sender joined. |
| Receiver Join Time | Specifies the time at which the receiver joined. |
| **Fields Specific for IPFM Mode** | |
| Priority | Specifies the flow priority for flows. |
| Policed | Specifies whether a flow is policed or not policed. |
| Receiver | Specifies the IP Address or the Host alias of the receiver joining the group. |
| Bandwidth | Specifies the bandwidth that is allotted for the traffic. |
| QOS/DSCP | Specifies the Switch-defined QoS Policy. |
| Policy ID | Specifies the policy ID applied to the multicast IP. |

| Field Specific for Generic Multicast Mode | |
|---|---|
| Receiver Interface IP | Specifies the IP address of the receiver interface joining the group. |

The following table describes the fields that appear on the **Inactive** tab.

*Table 20: Inactive Tab Fields and Descriptions*

| Field | Description |
|---|---|
| **Common Fields for IPFM and Generic Multicast Modes** | |
| VRF | Specifies the name of the VRF for the flow. |
| Multicast IP | Specifies the multicast IP address for the flow.<br>**Note**<br>You can click the chart link next to the Multicast IP address to view the pictorial representation of flow statistics. |
| Flow Alias | Specifies the name of the Flow Alias. |
| NAT | Specifies whether the flow is ingress, egress, or both ingress and egress. |
| Sender | Specifies the IP Address or the Host alias of the sender for the multicast group. |
| Sender Start Time | Displays the time from when the sender joined. |
| Receiver Join Time | Specifies the time at which the receiver joined. |
| **Fields Specific for IPFM Mode** | |
| Priority | Specifies the flow priority for flows. |
| Policed | Specifies whether a flow is policed or not policed. |
| Receiver | Specifies the IP Address or the Host alias of the receiver joining the group. |
| Bandwidth | Specifies the bandwidth that is allotted for the traffic. |
| QOS/DSCP | Specifies the Switch-defined QoS Policy. |
| Policy ID | Specifies the policy ID applied to the multicast IP. |
| Fault Reason | Specifies reason for the inactive flow.<br>Cisco Nexus Dashboard Fabric Controller determines the inactive flow if both the sender and receiver mroute exists with any of the following combinations.<br>• Receiver IIF is null<br>• Receiver OIF is null<br>• Sender IIF is null<br>• Sender OIF is null<br>In this scenario, the switch will not have any fault reason. Therefore, there is no fault reason for such inactive flows. |
| **Field Specific for Generic Multicast Mode** | |

| | |
|---|---|
| Receiver Interface IP | Specifies the IP address of the receiver interface joining the group. |

The following table describes the fields that appear on the **Sender Only** tab.

***Table 21: Sender Only Tab Field and Description***

| Field | Description |
|---|---|
| **Common Fields for IPFM and Generic Multicast Modes** | |
| VRF | Specifies the name of the VRF for the flow. |
| Multicast IP | Specifies the multicast IP address for the flow. |
| Flow Alias | Specifies the name of the Flow Alias. |
| Flow Link State | Specifies the flow link state, if it's allow or deny.<br>Click the **senderonly** link to view the network diagram or topology of the Sender and Receiver.<br>The dotted line displays the direction of the flow of traffic. You can hover over the nodes to view information. The table on the right side shows information about the Sender and Receiver.<br>The flows in the network diagram or topology show the multicast IP as well as the VRF. If the VRF is **default**, then the VRF will not be shown along with the multicast IP. |
| Sender | Specifies the name of the sender. |
| NAT | Specifies whether the flow is ingress, egress, or both ingress and egress. |
| Sender Switch | Specifies the IP address of the sender switch. |
| Sender Ingress Interface | Specifies the name of the sender ingress interface. |
| Sender Start Time | Displays the time from when the sender switch is transmitting information. |
| **Fields Specific for IPFM Mode** | |
| Policed | Specifies whether a flow is policed or not policed. |
| Policy ID | Specifies the policy ID applied to the multicast IP. |
| Bandwidth | Specifies the bandwidth that is allotted for the traffic. |
| QOS/DSCP | Specifies the Switch-defined QoS Policy. |
| Priority | Specifies the flow priority for flows. |

The following table describes the fields that appear on the **Receiver Only** tab.

***Table 22: Receiver Only Tab Field and Description***

| Field | Description |
|---|---|
| **Common Fields for IPFM and Generic Multicast Modes** | |
| VRF | Specifies the name of the VRF for the flow. |
| Multicast IP | Specifies the multicast IP address for the flow. |

| Field | Description |
|---|---|
| **Common Fields for IPFM and Generic Multicast Modes** | |
| Flow Alias | Specifies the name of the Flow Alias. |
| Flow Link State | Specifies the flow link state, if it's allow or deny.<br>Click the **receiveronly** link to view the network diagram or topology of the Sender and Receiver.<br>The dotted line displays the direction of the flow of traffic. You can hover over the nodes to view information. The table on the right side shows information about the Sender and Receiver.<br>The flows in the network diagram or topology show the multicast IP as well as the VRF. If the VRF is **default**, then the VRF will not be shown along with the multicast IP. |
| Source Specific Sender | Specifies the IP address of the multicast sender. |
| Receiver | Specifies the receiver ID. If the multicast receiver is remote, the **Remote** label can be seen next to its name. |
| NAT | Specifies whether the flow is ingress, egress, or both ingress and egress. |
| Receiver Switch | Specifies the IP address of the receiver switch. |
| Receiver Interface | Specifies the name of the destination switch interface. |
| Receiver Join Time | Specifies the time at which the receiver joined. |
| **Fields Specific for IPFM Mode** | |
| Bandwidth | Specifies the bandwidth that is allotted for the traffic. |
| Policy ID | Specifies the policy ID applied to the multicast IP. |
| Priority | Specifies the flow priority for flows. |
| QOS/DSCP | Specifies the Switch-defined QoS Policy. |



**Note** If stats are enabled on switches, only then they can be seen in Nexus Dashboard Fabric Controller.

Click the **Show** drop-down list in the statistical representation area to display the statistical data in various formats.

Click the arrow to export the statistical data. You can export it in `.csv` or `.pdf` formats.

**Note** Cisco Nexus Dashboard Fabric Controller holds the Flow statistics values in the Nexus Dashboard Fabric Controller server internal memory. Therefore, after a Nexus Dashboard Fabric Controller Restart or HA switch over, the Flow statistics won't show previously collected values. However, you can see the Flow statistics that are collected after the server Restart or HA switch over.

If the new flow joins before the uplinks between the switches that are detected in Nexus Dashboard Fabric Controller, a message BW_UNAVAIL appears. This is resolved after the uplinks between the switches are detected by Nexus Dashboard Fabric Controller after discovery of the devices.

# Flow Policies

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Hosts** > **Flow Policies**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Hosts** > **Flow Policies**.

Use this window to configure the flow policies.

**Note** When a user logs in to Nexus Dashboard Fabric Controller with a network operator role, all the buttons or options to add, delete, modify, import, export, or deploy policies are disabled. This user can only view policies, and deployment status or history.

The default policies are displayed on the **Flow Policies** tab. By default, the bandwidth of these policies is 0. You can configure the bandwidth such that any flow that matches the default flow policy will accordingly use the bandwidth and QOS/DSCP parameters. The policy is deployed to all the devices when you save the configuration.

**Note** When you undeploy a default policy, it will be reset to default values, that is, Bandwidth:0gbps, DSCP:Best Effort, and Policer:Enabled.

Policies are automatically deployed to switches whenever they are created, edited, or imported. You can choose to undeploy or redeploy the policies, by choosing the appropriate actions in the **Actions** drop-down list. The policies will not be deployed correctly if the device is rebooting while the policy was deployed. In such case, the **Failed** message appears in the **Deployment Status** column.

The default flow policies must be deployed successfully to the switch before you deploy the custom flow policies on that switch. Otherwise, the custom policies will fail to deploy. Ensure that you deploy all default policies successfully to all the switches before you add, edit, import, or deploy custom policies.

**Note** If you have created a custom or non-default VRF, although the host and flow policies are automatically created for the VRF, use the action options in this window to manually deploy the flow policies to the switches in the fabric.

The following table describes the fields that appear on this page.

*Table 23: Flow Policies Table Field and Description*

| Field | Description |
|---|---|
| VRF | Specifies the name of the VRF for the flow policy. |
| Policy Name | Specifies the flow policy name. |
| Multicast IP Range | Specifies the multicast IP address for the traffic. Click **view** to view the details such as starting and ending IP addresses of the multicast range as well as the flow priority in the **Multicast Range List** box. |
| Bandwidth | Specifies the bandwidth that is allotted for the traffic. |
| QoS/DSCP | Specifies the Switch-defined QoS Policy. |
| Deployment Action | Specifies the action that is performed on the switch for that host policy.<br>• **Create** - The policy is deployed on the switch.<br>• **Delete** - The policy is undeployed from the switch. |
| Deployment Status | Specifies if the flow policy is deployed successfully, not deployed, or failed. |
| In Use | Specifies if the flow policy is in use or not. |
| Policer | Specifies whether the policer for a flow policy is enabled or disabled.<br>**Note**<br>In adding or editing a flow policy, the default policer state is **Enabled**. |
| Last Updated | Specifies the date and time at which the flow policy was last updated.<br>The format is *Day MMM DD YYYY HH:MM:SS Timezone*. |

Click the table header to sort the entries in alphabetical order of that parameter.

The following table describes the action items, in the **Actions** drop-down list, that appears in the **Flow Policies** horizontal tab on the **Flows** tab in the **Fabric Overview** window.

**Note** A new flow policy or an edited flow policy is effective only under the following circumstances:

- If the new flow matches the existing flow policy.
- If the flow expires and reforms, while the new policy is already created or edited, that matches with the flow policy.

*Table 24: Flow Policies Actions and Description*

| Field | Description |
|---|---|
| Create Flow Policy | Allows you to create a new flow policy. For more information, see Creating a Flow Policy, on page 247. |
| Edit Flow Policy | Allows you to view or edit the selected flow policy parameters.<br>**Note**<br>The default flow policies must be deployed successfully to the switch before you deploy the custom flow policies on that switch. Otherwise, the custom policies will fail to deploy. Ensure that you deploy all default policies successfully to all the switches before you edit custom policies.<br><br>To edit a flow policy for a VRF, select the check box next to the VRF and choose **Edit Flow Policy** action. In the **Edit Flow Policy** window, you can make the required changes and click **Save & Deploy** to deploy the changes or click **Cancel** to discard the changes.<br><br>The deployment completed message appears at the bottom of the window. You can click **Refresh** to refresh the current deployment status in the window or click **View Details** to verify the deployment details. |
| Delete Flow Policy | Allows you to delete the user-defined flow policy.<br>**Note**<br>• You cannot delete the default flow policies.<br>• Undeploy policies from all switches before deleting them from Nexus Dashboard Fabric Controller.<br>• You can select more than one flow policy to delete.<br><br>To delete a flow policy, select the check box next to that VRF and choose the **Delete Flow Policy** action. A warning message appears asking you to undeploy policies from the switches. Click **Confirm** to proceed with deletion and leave the policies on the switches or click **Cancel** to discard the delete operation. |

| Field | Description |
|---|---|
| Purge | Allows you to delete all the flow policies at a single instance.<br>**Note**<br>Undeploy policies from all switches before deleting them from Nexus Dashboard Fabric Controller.<br>To delete all flow policies, choose the **Purge** action. A warning message appears asking you to undeploy policies from all the switches. Click **Confirm** to proceed with deletion and leave the policies on the switches or click **Cancel** to discard the delete operation. |
| Import | Allows you to import flow policies from a csv file.<br>**Note**<br>The default flow policies must be deployed successfully to the switch before you deploy the custom flow policies on that switch. Otherwise, the custom policies will fail to deploy. Ensure that you deploy all default policies successfully to all the switches before you import custom policies.<br>After import, all policies imported from a csv file are applied to all managed switches automatically.<br>To import the flow policies, choose the **Import** action. Browse the directory and select the .csv file that contains the flow policy configuration information. The policy will not be imported if the format in the .csv file is incorrect. Click **Open**. The imported policies are automatically deployed to all the switches in the fabric. |
| Export | Allows you to export flow policies to a `csv` file.<br>To export the flow policies, choose the **Export** action. Select a location on your local system directory to store the flow policy details file. Click **Save**. The flow policy file is exported to your local directory. The file name is appended with the date on which the file is exported. The format of the exported file is .csv. |
| Deploy Selected Policies | Select this option to deploy only the selected policies to the devices. You can deploy other policies when required.<br>Select one or more check boxes next to the policy name. Select this option to deploy selected policies to the switch. |
| Deploy All Custom Policies | Select this option to deploy all the custom or user-defined policies at a single instance.<br>The policies are deployed even if the switch is rebooting. In such case, the deployment fails and a status message **Failed** appears in the **Deployment Status** column. |
| Deploy All Default Policies | Select this option to deploy all default policies to the switch. |
| Undeploy Selected Policies | Select this option to undeploy the selected policies.<br>To undeploy the selected policies, select one or more check boxes next to the VRFs. Select this option from the drop-down list to undeploy the selected policies. |
| Undeploy All Custom Policies | Select this option to undeploy all the custom or user-defined policies at a single instance. |

| Field | Description |
|---|---|
| Undeploy All Default Policies | Select this option to undeploy all the default policies at a single instance. |
| Redo All Failed Policies | The deployment or undeployment of policies may fail due to various reasons. Select this option to deploy all the failed policies.<br>All the deployments that failed on switches previously will be deployed again to only those switches. All the undeployments that failed on switches previously will be undeployed again from only those switches. |
| Deployment History | Select this option to view the deployment history of the selected policy for the switch in the **Deployment History** pane.<br>The **Deployment History** pane displays the following fields:<br>• Policy Name - Specifies the selected policy name.<br>• VRF - Specifies the VRF for the selected policy.<br>• Switch Name - Specifies the name of the switch that the policy was deployed to.<br>• Deployment Status - Displays the status of deployment. It shows if the deployment was a success, failed, or not deployed. Click on the deployment status, for example, **Success**, to see more details. For more information about the deployment status, see Deployment Status, on page 246.<br>• Action - Specifies the action that is performed on the switch for that flow policy.<br>  • **Create** - Implies that the policy has been deployed on the switch.<br>  • **Delete** - Implies that the policy has been undeployed from the switch.<br>• Deployment Date/Time - Specifies the date and time at which the host policy was last updated. The format is Day MMM DD YYYY HH:MM:SS Timezone .<br>• Failed Reason - Species why the policy was not successfully deployed. |

### Deployment Status

The following table describes the fields that appear on the Deployment Status.

*Table 25: Deployment Status Field and Description*

| Field | Description |
|---|---|
| Policy Name | Specifies the name of the flow policy. |
| VRF | Specifies the name of the VRF. |
| Switch Name | Specifies the switch on which the VRF is deployed. |
| IP Address | Specifies the IP address of the switch. |

| Field | Description |
|---|---|
| Deployment Status | Displays the status of the deployment. It shows if the deployment was a **Success** or **Failed** along with the reason why the deployment failed. |
| Action | Specifies the action that is performed on the switch, for example, **Create**. |
| Deployment Date/Time | Displays the date and time when the deployment was initialized. |

This section contains the following:

## Creating a Flow Policy



**Note** The default host policies must be deployed successfully to the switch before you deploy the custom host policies on that switch. Otherwise, the custom policies will fail to deploy. Ensure that you deploy all the default policies successfully to all the switches before you add custom policies.

To create a flow policy from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Procedure

**Step 1** Click **Actions** and choose **Create Flow Policy**.

The **Create Flow Policy** window is displayed.

**Step 2** In the **Create Flow Policy** window, specify the parameters in the following fields.

- **VRF** - Click the **Select a VRF** link to open the **Select a VRF** window. The default VRF is also listed in the window. Search and select a VRF for the host and click **Save**.

  **Note**
  - Policy names can be repeated across VRFs, that is, they are unique only within a VRF.
  - Across the VRF, host policies may be same or different.
  - Sequence number for the host policies is per VRF.

- **Policy Name** - Specify a unique policy name for the flow policy.
- **Bandwidth** - Specifies the bandwidth that is allocated for the flow policy. Select of the radio buttons to choose **Gbps**, **Mbps**, or **Kbps**.

**Step 3** From the **QoS/DSCP** drop-down list, choose an appropriate ENUM value.

**Step 4** Click the **Policer** check box to enable or disable policer for a flow.

**Step 5** In **Multicast IP Range**, enter the beginning IP and ending IP Address for the multicast range in the **From** and **To** fields. The valid range is between 224.0.0.0 and 239.255.255.255.

From the **Flow Priority** drop-down list, choose the priority for the flow. You can choose either **Default** or **Critical**. The default value is **Default**.

The flow priority is used during the following scenarios:

- Error Recovery - Unicast Routing Information Base (URIB) reachability changes on flows, and a re-Reverse-path forwarding (RPF) is being performed. When a set of existing flows is retried, the recovery starts from the flows with **Critical** priority.
- Flow Retry - When pending flows are retried, the **Critical** priority flows are retried first.

**Actions** - Actions has a variety of icons to perform various actions. Click the tick mark icon if you have entered the correct details; if not, click the check mark icon to add the multicast range to the policy. Click the edit icon if you want to modify the details or click the bin icon to delete the row. Click the Plus (+) mark to add another row.

**Step 6** Click **Save & Deploy** to deploy the new policy or click **Cancel** to discard the changes. The deployment completed message appears at the bottom of the window. You can click **Refresh** to refresh the current deployment status in the window or click **View Details** to verify the deployment details.

## Flow Alias

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Flows** > **Flow Alias**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Flows** > **Flow Alias**.

Use this tab to configure flow alias.



**Note** This section is applicable for both the IPFM and Generic Multicast modes in Nexus Dashboard Fabric Controller.

Using the Flow Alias feature, you can specify names for multicast groups. The multicast IP addresses are difficult to remember, thus by assigning a name to the multicast IP address, you can search and add policies based on the name.

The following table describes the fields that appear in this window.

***Table 26: Flow Alias Table Field and Description***

| Field | Description |
|---|---|
| VRF | Specifies the VRF for the flow alias. |
| Policy Name | Specifies the policy name. |
| Multicast IP Range | Specifies the multicast IP address for the traffic. |

| Field | Description |
|---|---|
| Description | Description added to the flow alias. |
| Last Updated | Specifies the date on which the flow alias was last updated. |

The following table describes the action items, in the **Actions** drop-down list, that appears in the **Flow Alias** horizontal tab on the **Flows** tab of the **Fabric Overview** window.

*Table 27: Flow Alias Actions and Description*

| Action Item | Description |
|---|---|
| Create Flow Alias | Allows you to create a new flow alias. For instructions about creating a new flow alias, see Creating Flow Alias, on page 249. |
| Edit Flow Alias | Allows you to view or edit the selected flow alias parameters.<br>To edit the flow alias, select the check box next to the flow alias that you want to delete and choose **Edit Flow Alias**. In the **Edit Flow Alias** window, edit the required values and click **Submit** to apply the changes or click **Cancel** to discard the flow alias. The edited flow alias is shown in the table in the **Flow Alias** window. |
| Delete Flow Alias | Allows you to delete the flow alias.<br>To delete a flow alias, select the check box next to the flow alias that you want to delete and choose **Delete Flow Alias**. You can select multiple flow alias entries and delete them at the same instance. |
| Import | Allows you to import flow aliases for devices in the fabric.<br>To import flow aliases, choose **Import**. Browse the directory and select the `.csv` file that contains the flow IP address and corresponding unique flow name information. Click **Open**. The flow aliases are imported and displayed in the **Flow Alias** window. |
| Export | Allows you to export flow aliases for devices in the fabric.<br>To export a flow alias, choose **Export**. Select a location on your local system directory to store the flow aliases configuration from Nexus Dashboard Fabric Controller and click **Save**. The flow alias configuration file is exported to your local directory. The file name is appended with the date and time at which the file was exported. The format of the exported file is `.csv`. |

This section contains the following:

## Creating Flow Alias

### UI Navigation

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Flows** > **Flow Alias**.

- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Flows** > **Flow Alias**.

To create a flow alias from the Cisco Nexus Dashboard Fabric Controller, perform the following steps:

### Procedure

**Step 1** In the **Flow Alias** window, from the **Actions** drop-down list, choose **Create Flow Alias**.

**Step 2** In the **Create Flow Alias** window, enter the following:

**Note**
All the fields are mandatory.

- **VRF** - Select the VRF from this drop-down list. The default value is **default**.

  **Note**
  Host and IP Address are unique per VRF, that is, same host name with the same IP Address can exist in multiple VRFs.

- **Flow Name** - Enter a fully qualified unique flow name for identification of the flow alias.
- **Multicast IP Address** - Enter the multicast IP address for the flow alias.
- **Description** - Enter a description for the flow alias.

**Step 3** Click **Submit** to apply the changes.

Click **Cancel** to discard the flow alias.

The new flow alias is shown in the table in the **Flow Alias** window.

## Static Flow

### UI Navigation

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Hosts** > **Static Flow**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Hosts** > **Static Flow**.

You configure a static receiver using the **Static Flow** window. Use the **Select an Option** field to select a switch before creating a static flow for it.

***Table 28: Static Flow Actions and Description***

| Field | Description |
|---|---|
| Create Static Flow | Allows you to create a static flow. For more information, see Creating a Static Flow, on page 251. |

| Field | Description |
|---|---|
| Delete Static Flow | Allows you to delete the static flow.<br>Select a static flow that you need to delete and click the **Delete Static Flow** action to delete the selected static flow. |

*Table 29: Static Flow Table Field and Description*

| Field | Description |
|---|---|
| VRF | Specifies the VRF for a static flow. |
| Group | Specifies the group for a static flow. |
| Source | Specifies the source IP address for the static flow. |
| Interface Name | Specifies the interface name for the static flow. If it is not specified while creating the static flow, it is displayed as **N/A**. |
| Deployment Action | Specifies the action that is performed on the switch for the rule. Create implies that the static flow has been deployed on the switch. Delete implies that the static flow has been undeployed from the switch. |
| Deployment Status | Specifies if the static flow is deployed or not. If there is a deployment failure, hover over the information icon to view the failure reason. |
| Last Updated | Specifies the date and time at which the static flow was last updated.<br>The format is Day MMM DD YYYY HH:MM:SS Timezone. |

## Creating a Static Flow

To create a static flow for the selected switch, perform the following steps:

### Before you begin

Select a switch in the **Static Flow** tab of the **Fabric Overview** window before creating a static flow for it.

### Procedure

**Step 1** Click **Actions** and choose **Create Static Flow**.

The **Create Static Flow** window is displayed.

**Step 2** In the **Create Static Flow** window, specify the parameters in the following fields.

**Switch** - Specifies the switch name. This field is read-only, and it is based on the switch selected in the **Static Flow** window.

**Group** - Specifies the multicast group.

**Source** - Specifies the source IP address.

**Interface Name** - Specify the interface name for the static flow. This field is optional. If you do not specify an interface name, the host IP 0.0.0.0 is passed to the API and config is created using Null0 interface.

**Step 3** Click **Save & Deploy** to save the static flow.

Click **Cancel** to discard it.

# Metrics

The Metric tab displays the infrastructure health and status. You can view CPU utilization, Memory utilization, Traffic, Temperature, Interface, and Links details.

The following table describes the columns that appears on **CPU** and **Memory** tab.

| Fields | Descriptions |
|---|---|
| Switch Name | Specifies the name of switch. |
| IP Address | Specifies the switch IP address. |
| Low Value (%) | Specifies the lowest CPU utilization value on the switch. |
| Avg. Value (%) | Specifies the average CPU utilization value on the switch. |
| High Value (%) | Specifies the high CPU utilization value on the switch. |
| Range Preview | Specifies the linear range preview. |
| Last Update Time | Specifies the last updated time on the switch. |
| Show last day | Click **Show last day** to view data for selected day, week, month, and year. |

The following table describes the columns that appears on **Traffic** tab.

| Fields | Descriptions |
|---|---|
| Switch Name | Specifies the name of switch. |
| Avg. Rx | Specifies the average Rx value. |
| Peak Rx | Specifies the peak Rx value. |
| Avg. Tx | Specifies the average Tx value. |
| Peak Tx | Specifies the peak Tx value. |
| Avg. Rx+Tx | Specifies the average of Rx and Tx value. |
| Avg. Errors | Specifies the average error value. |
| Peak Errors | Specifies the peak error value. |
| Avg. Discards | Specifies the average discard value. |
| Peak Discards | Specifies the peak discard value. |
| Last Update Time | Specifies the last updated time. |
| Show last day | Click **Show last day** to view data for selected day, week, month, and year. |

The following table describes the columns that appears on **Temperature** tab.

| Fields | Descriptions |
|---|---|
| Switch Name | Specifies the name of switch. |
| IP Address | Specifies the switch IP address. |
| Temperature Module | Specifies the module of temperature. |
| Low Value (C) | Specifies the lowest temperature value. |
| Avg. Value (C) | Specifies the average temperature value. |
| High Value (C) | Specifies the high temperature value. |
| Show last day | Click **Show last day** to view data for selected day, week, month, and year. |

The following table describes the columns that appears on **Interface** tab.

| Fields | Descriptions |
|---|---|
| Switch | Specifies the name of switch. |
| Interface | Specifies the name of interface |
| Description | Specifies the description of interface. |
| Speed | Specifies the speed of the interface. |
| Status | Specifies the status of switch link. |
| **Rx.** | |
| Avg. | Specifies the average Rx value. |
| Avg% | Specifies the average percentage of Rx value. |
| Peak | Specifies the peak Rx value. |
| Peak% | Specifies the peak percentage Rx value. |
| **Tx.** | |
| Avg. | Specifies the average Tx value. |
| Avg% | Specifies the average percentage of Tx value. |
| Peak | Specifies the peak Tx value. |
| Peak% | Specifies the peak percentage Tx value. |
| Rx+Tx | Specifies the sum value of Rx and Tx. |
| **Errors** | |
| In Avg. | Specifies the in average error value. |
| Out Avg. | Specifies the out peak error value. |
| In Peak | Specifies the in peak error value. |
| Out Peak | Specifies the out peak error value. |

| Fields | Descriptions |
|---|---|
| **Discards** | |
| In Avg. | Specifies the average discard value. |
| Out Avg. | Specifies the peak discard value. |
| In Peak | Specifies the in peak discard value. |
| Out Peak | Specifies the out peak discard value. |
| **Show last day** | Click **Show last day** to view data for selected day, week, month, and year. |

The following table describes the columns that appears on **Link** tab.

| Fields | Descriptions |
|---|---|
| Switch | Specifies the name of switch. |
| Vlans | Specifies the VLAN name. |
| Speed | Displays the speed value. |
| Status | Specifies the status of switch. |
| **Rx.** | |
| Avg. | Specifies the average Rx value. |
| Avg% | Specifies the average percentage of Rx value. |
| Peak | Specifies the peak Rx value. |
| Peak% | Specifies the peak percentage Rx value. |
| **Tx.** | |
| Avg. | Specifies the average Tx value. |
| Avg% | Specifies the average percentage of Tx value. |
| Peak | Specifies the peak Tx value. |
| Peak% | Specifies the peak percentage Tx value. |
| Rx+Tx | Specifies the sum value of Rx and Tx. |
| **Errors** | |
| In Avg. | Specifies the in average error value. |
| Out Avg. | Specifies the out peak error value. |
| In Peak | Specifies the in peak error value. |
| Out Peak | Specifies the out peak error value. |
| **Discards** | |
| In Avg. | Specifies the average discard value. |

| Fields | Descriptions |
|---|---|
| Out Avg. | Specifies the peak discard value. |
| In Peak | Specifies the in peak discard value. |
| Out Peak | Specifies the out peak discard value. |
| **Show last day** | Click **Show last day** to view data for selected day, week, month, and year. |

# Multicast NAT

Multicast NAT translation of UDP stream is supported on the Nexus Dashboard Fabric Controller IPFM mode. You can apply NAT for the incoming traffic (ingress), or on the egress link or interface. The scope of ingress NAT is entire switch, whereas egress NAT is for a specific interface. The same switch can have both ingress and egress NAT. However, it can't be on the same flow for a given switch. Egress NAT has capability to replicate the same flow up to 40 times. To achieve this function, the service-reflect interface is defined on the switch. It serves for multiple or single egress port.



**Note** Ingress and/or Egress NAT translation is supported only on the sender switch, also known as First Hop Router (FHR), and receiver switch, also known as Last Hop Router (LHR). It is not supported on intermediates nodes such as spine switches.

For more information about NAT, see *Cisco Nexus 9000 Series NX-OS IP Fabric for Media Solution Guide*.

### Prerequisites

- Set up loopback interface with PIM sparse mode. When flow is translated, post-translated source needs to be secondary IP address on this loopback to make sure RPF check won't fail. This loopback is configured as service reflect interface for NAT purpose. You need to set up lookback per VRF.

  Here is an example to configure the loopback interface:

```
interface loopback10
ip router ospf 1 area 0
ip pim sparse-mode
ip address 192.168.1.1/32
ip address 172.16.1.10/32 secondary

ip service-reflect source-interface loopback10
```

- TCAM memory carving must be completed.

  The command to configure the TCAM for Multicast NAT is as follows:

```
hardware access-list tcam region mcast-nat tcam-size
```

  For information about switch models that support multicast NAT, see Configuring Multicast Service Reflection with NBM in Cisco Nexus 9000 Series NX-OS IP Fabric for Media Solution Guide.

# NAT Modes

NAT Mode objects are created per switch and VRF. The switches are populated in the drop-down based on the scope. You should select the switch to list and operate on the corresponding NAT Mode objects.

Choose **LAN > Fabrics**. Double-click a fabric name and click **Multicast NAT > NAT Modes** to configure NAT modes.

The following table describes the fields that appear on the **NAT Modes** tab.

| Field | Description |
|---|---|
| VRF | Specifies the VRF for the multicast NAT. VRF support is not applicable for eNAT, however, it is applicable for iNAT. |
| Group | Specifies the multicast address of the NAT mode. |
| Mode | Specifies the multicast NAT mode, that is, ingress or egress. |
| Deployment Action | Specifies the action that is performed on the switch for that mode. Create implies that the mode has been deployed on the switch. Delete implies that the mode has been undeployed from the switch. |
| Deployment Status | Specifies if the mode is deployed or not. If there's deployment failure, hover over the information icon to view the failure reason. |
| Last Updated | Specifies the date and time at which the mode was last updated.<br>The format is Day MMM DD YYYY HH:MM:SS Timezone. |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on the **NAT Modes** tab.

| Action Item | Description |
|---|---|
| Create NAT Mode | Choose **Create NAT Mode** to add a NAT mode. |
| Delete NAT Mode | Select a mode from the table and choose **Delete NAT Mode** to delete the mode. |
| Import | Allows you to import NAT modes from a CSV file to Nexus Dashboard Fabric Controller. |
| Export | Allows you to export NAT modes from Nexus Dashboard Fabric Controller to a CSV file. |
| Deploy Selected NAT Modes | Select modes from the table and choose **Deploy Selected NAT Modes** to deploy selected modes to the switch. |
| Deploy All NAT Modes | Choose **Deploy All NAT Modes** to deploy all modes to the switch. |
| Undeploy Selected NAT Modes | Select modes from the table and choose **Undeploy Selected NAT Modes** to undeploy selected modes from the switch. |

| Action Item | Description |
|---|---|
| Undeploy All NAT Modes | Choose **Undeploy All NAT Modes** to undeploy all modes from the switch. |
| Redo All Failed NAT Modes | Choose **Redo All Failed NAT Modes** to deploy all failed modes. |
| Deployment History | Select a mode from the table and choose **Deployment History** to view the deployment history of the selected mode.<br>Deployment History shows the following fields.<br>• Switch Name—Specifies the name of the switch that the mode was deployed to.<br>• VRF—Specifies the name of the VRF that mode was deployed to.<br>• Group—Specifies the multicast group of the NAT mode.<br>• Mode—Specifies the NAT mode, that is, ingress or egress.<br>• Deployment Status—Displays the status of deployment. It shows if the deployment was Success or Failed.<br>• Action—Specifies the action that is performed on the switch for that mode. Create implies that the mode has been deployed on the switch. Delete implies that the mode has been undeployed from the switch.<br>• Deployment Date/Time—Specifies the date and time at which the mode was last updated. The format is Day MMM DD YYYY HH:MM:SS Timezone.<br>• Failed Reason—Specifies why the mode wasn't successfully deployed. |

## Adding a NAT Mode

### Procedure

**Step 1** Choose **LAN** > **Fabrics**.

**Step 2** Double-click a fabric name.

The **Fabric Overview** window appears.

**Step 3** Click the **Multicast NAT** tab.

**Step 4** Click the **NAT Modes** tab.

**Step 5** Click **Actions > Create NAT Mode** to add a NAT mode.

The **Add NAT Mode** window appears.

**Step 6** In the **Add NAT Mode** window, specify the following information:

**Mode**: Select the multicast NAT mode, that is, **Ingress** or **Egress**.

**Selected Switch**: Specifies the switch name. This field is read-only, and it's based on the switch selected in the **NAT Modes** tab.

**VRF**: Select the VRF to which the NAT mode should belong to. For the **Egress** NAT mode, the default VRF is selected and it's non-editable.

**Group / Mask**: Specify the multicast group with the mask. The same group can't be ingress as well as egress NAT on a given switch. You need to identify whether particular group or mask would be ingress or egress.

**Step 7** Click **Save & Deploy** to save the NAT mode and deploy it.

### Deleting a NAT Mode

#### Procedure

**Step 1** Choose **LAN > Fabrics**.

**Step 2** Double-click a fabric name.

The **Fabric Overview** window appears.

**Step 3** Click the **Multicast NAT** tab.

**Step 4** Click the **NAT Modes** tab.

**Step 5** Select the NAT mode that you need to delete and click **Actions > Delete NAT Mode** to delete a NAT mode.

If the NAT mode isn't deployed or failed, you can skip this step.

**Step 6** Click **Confirm** to delete the selected NAT mode.

## Egress Interface Mappings

Choose **LAN > Fabrics**. Double-click a fabric name and click **Multicast NAT > Egress Interface Mappings** to configure egress interface mappings.

The following table describes the fields that appear on the **Egress Interface Mappings** tab.

| Field | Description |
|---|---|
| Egress Interfaces | Specifies the egress interfaces for the mapping. |

| | |
|---|---|
| Map Interface | Specifies the map interface.<br>Egress interfaces and map interface have Many to One relationship. When there are more than one Egress Interfaces for a mapping, it is shown as a hyperlink. You can click on the hyperlink to see the complete list of interfaces. |
| Max Replications | Specifies the max replications for the map interface. |
| Deployment Action | Specifies the action that is performed on the switch for that egress interface mapping. Create implies that the egress interface mapping has been deployed on the switch. Delete implies that the egress interface mapping has been undeployed from the switch. |
| Deployment Status | Specifies if the egress interface mapping is deployed or not. If there's deployment failure, hover over the information icon to view the failure reason. |
| Last Updated | Specifies the date and time at which the egress interface mapping was last updated.<br>The format is Day MMM DD YYYY HH:MM:SS Timezone. |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on the **Egress Interface Mappings** tab.

| Action Item | Description |
|---|---|
| Create NAT Egress Interface Mapping | Choose **Create NAT Egress Interface Mapping** to add an egress interface mapping. |
| Edit NAT Egress Interface Mapping | Select a mode from the table and choose **Edit NAT Egress Interface Mapping** to edit an egress interface mapping. |
| Delete NAT Egress Interface Mapping | Select a mode from the table and choose **Delete NAT Egress Interface Mapping** to delete an egress interface mapping. |
| Import | Allows you to import NAT egress interface mappings from a CSV file to Nexus Dashboard Fabric Controller. |
| Export | Allows you to export NAT egress interface mappings from Nexus Dashboard Fabric Controller to a CSV file. |
| Deploy Selected NAT Egress Interface Mappings | Select modes from the table and choose **Deploy Selected NAT Egress Interface Mappings** to deploy selected egress interface mapping to the switch. |
| Deploy All NAT Egress Interface Mappings | Choose **Deploy All NAT Egress Interface Mappings** to deploy all egress interface mappings to the switch. |
| Undeploy Selected NAT Egress Interface Mappings | Select modes from the table and choose **Undeploy Selected NAT Egress Interface Mappings** to undeploy selected egress interface mappings from the switch. |

| Action Item | Description |
|---|---|
| Undeploy All NAT Egress Interface Mappings | Choose **Undeploy All NAT Egress Interface Mappings** to undeploy all egress interface mapping from the switch. |
| Redo All Failed NAT Egress Interface Mappings | Choose **Redo All Failed NAT Egress Interface Mappings** to deploy all failed egress interface mappings. |
| Deployment History | Select a mode from the table and choose **Deployment History** to view the deployment history of the selected egress interface mapping.<br>Deployment History shows the following fields.<br>• Switch Name—Specifies the name of the switch that the mode was deployed to.<br>• Map Interface—Specifies the map interface for the egress interface mappings.<br>• Max Replications—Specifies the maximum replications for the egress interface mappings.<br>• Egress Interfaces—Specifies the name of the egress interface that the mapping is deployed to.<br>• Deployment Status—Displays the status of deployment. It shows if the deployment was Success or Failed.<br>• Action—Specifies the action that is performed on the switch for that egress interface mapping. Create implies that the mapping has been deployed on the switch. Delete implies that the mapping has been undeployed from the switch.<br>• Deployment Date/Time—Specifies the date and time at which the mapping was last updated. The format is Day MMM DD YYYY HH:MM:SS Timezone.<br>• Failed Reason—Specifies why the mode wasn't successfully deployed. |

## Adding NAT Egress Interface Mapping

### Procedure

**Step 1** Choose **LAN > Fabrics**.

**Step 2** Double-click a fabric name.

The **Fabric Overview** window appears.

**Step 3** Click the **Multicast NAT** tab.

**Step 4** Click the **Egress Interface Mappings** tab.

**Step 5** Click **Actions > Create NAT Egress Interface Mapping** to add an egress interface mapping.

The **Add Egress Interface Mappings** window appears.

**Step 6** In the **Add Egress Interface Mappings** window, specify the following information:

**Selected Switch**: Specifies the switch name. This field is read-only, and it's based on the switch selected in the **Egress Interface Mappings** window.

**Egress Interfaces**: Specifies the egress interface. You can select one or more egress interfaces. Egress Interfaces and Map interface are prepopulated based on the switch selected.

You can select multiple Egress Interfaces by selecting the **Select one or more** option and click the **Select** option to choose the interfaces. The Select window shows the interfaces that are available, that is, the interfaces that are already defined in other mappings are filtered out. To select all the interfaces, you can select **All**. When **All** is selected, the option to select individual egress interfaces is disabled.

**Map Interface**: Specifies the map interface. An interface can either be an Egress Interface or a Map Interface and can't be both. An error is displayed if you select a map interface that is already selected as an Egress Interface.

**Max Replications**: Specifies the maximum replications for the map interface. The range for this field is 1–40. The default value is 40.

**Step 7** Click **Save & Deploy** to save the NAT mode and deploy it.

## Editing NAT Egress Interface Mapping

### Procedure

**Step 1** Choose **LAN > Fabrics**.

**Step 2** Double-click a fabric name.

The **Fabric Overview** window appears.

**Step 3** Click the **Multicast NAT** tab.

**Step 4** Click the **Egress Interface Mappings** tab.

**Step 5** Click **Actions > Edit NAT Egress Interface Mapping** to edit an egress interface mapping.

The **Edit Egress Interface Mappings** window appears.

**Step 6** In the **Edit Egress Interface Mappings** window, specify the following information:

Edit egress interfaces and **Max Replications** field. Specify the new value in **Max Replications** that should be within 1–40.

**Step 7** Click **Save & Deploy** to save the egress interface mapping and deploy it.

### Deleting Egress Interface Mapping

Deleting an egress interface mapping doesn't undeploy the egress interface mapping from the switch. Therefore, make sure to undeploy the egress interface mapping from the switch before deleting it from Nexus Dashboard Fabric Controller.

#### Procedure

**Step 1** Choose **LAN > Fabrics**.

**Step 2** Double-click a fabric name.

The **Fabric Overview** window appears.

**Step 3** Click the **Multicast NAT** tab.

**Step 4** Click the **Egress Interface Mappings** tab.

**Step 5** Click **Actions > Delete NAT Egress Interface Mapping** to delete the selected egress interface mapping.

**Step 6** Click **Confirm** to delete the selected egress interface mapping.

## NAT Rules

NAT rules are identical for ingress and egress NAT except you need to also specify receiver OIF for egress NAT.

Choose **LAN > Fabrics**. Double-click a fabric name and click **Multicast NAT > NAT Rules** to configure NAT rules.

The following table describes the fields that appear on the **NAT Rules** tab.

| Field | Description |
|---|---|
| VRF | Specifies the VRF for the multicast NAT. |
| Mode | Specifies the NAT mode, that is, ingress or egress. |
| Pre-Translation Group | Specifies the multicast group before NAT. |
| Post-Translation Group | Specifies the multicast group after NAT. |
| Group Mask | Specifies the group mask. |
| Pre-Translation Source | Specifies the source IP address before NAT. |
| Post-Translation Source | Specifies the source IP address after NAT. |
| Source Mask | Specifies the source mask. |
| Post-Translation Source Port | Specifies the source port after NAT. The range is 0–65535. The value 0 means that there's no translation of UDP source port. |
| Post-Translation Destination Port | Specifies the destination port after NAT. The value 0 means that there's no translation of UDP destination port. |
| Static Oif | Specifies the static outgoing interface to bind the Egress NAT rule to. This drop-down is populated with Egress Interfaces defined in the **Egress Interface Mappings** window. This field is disabled for Ingress mode. |

| | |
|---|---|
| Deployment Action | Specifies the action that is performed on the switch for the rule. Create implies that the rule has been deployed on the switch. Delete implies that the rule has been undeployed from the switch. |
| Deployment Status | Specifies if the rule is deployed or not. If there's a deployment failure, hover over the information icon to view the failure reason. |
| Last Updated | Specifies the date and time at which the rule was last updated.<br>The format is Day MMM DD YYYY HH:MM:SS Timezone. |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on the **NAT Rules** tab.

| Action Item | Description |
|---|---|
| Create NAT Rule | Choose **Create NAT Rule** to add a NAT rule. |
| Delete NAT Rule | Select a mode from the table and choose **Delete NAT Rule** to delete the rule. |
| Import | Allows you to import NAT rules from a CSV file to Nexus Dashboard Fabric Controller. |
| Export | Allows you to export NAT rules from Nexus Dashboard Fabric Controller to a CSV file. |
| Deploy Selected NAT Rules | Select rules from the table and choose **Deploy Selected NAT Rules** to deploy selected rules to the switch. |
| Deploy All NAT Rules | Choose **Deploy All NAT Rules** to deploy all rules to the switch. |
| Undeploy Selected NAT Rules | Select rules from the table and choose **Undeploy Selected NAT Rules** to undeploy selected rules to the switch. |
| Undeploy All NAT Rules | Choose **Undeploy All NAT Rules** to undeploy all rules from the switch. |
| Redo All Failed NAT Rules | Choose **Redo All Failed NAT Rules** to deploy all failed rules. |

| Action Item | Description |
|---|---|
| Deployment History | Select a rule from the table and choose **Deployment History** to view the deployment history of the selected rule.<br>Deployment History shows the following fields.<br>• Switch Name—Specifies the name of the switch that the rule was deployed to.<br>• VRF-Specifies the VRF that the mapping belongs to.<br>• Deployment Status—Displays the status of deployment. It shows if the deployment was Success or Failed.<br>• Action—Specifies the action that is performed on the switch for that rule. Create implies that the rule has been deployed on the switch. Delete implies that the rule has been undeployed from the switch.<br>• Deployment Date/Time—Specifies the date and time at which the rule was last updated. The format is Day MMM DD YYYY HH:MM:SS Timezone.<br>• Failed Reason — Specifies why the rule wasn't successfully deployed. |

## Adding NAT Rule

### Procedure

**Step 1** Choose **LAN > Fabrics**.

**Step 2** Double-click a fabric name.

The **Fabric Overview** window appears.

**Step 3** Click the **Multicast NAT** tab.

**Step 4** Click the **NAT Rules** tab.

**Step 5** Click **Actions > Create NAT Rule** to add a NAT rule.

The **Add NAT Rule** window appears.

**Step 6** In the **Add NAT Rule** window, specify the following information:

**Mode**: Select the NAT mode, that is, **Ingress** or **Egress**.

**Selected Switch**: Specifies the switch name. This field is read-only, and it's based on the switch selected in the **NAT Rules** tab.

**VRF**: Select the VRF for the NAT rule. By default, it's the **default** VRF.

**Pre-Translation Group**: Specifies the multicast group before NAT.

**Post-Translation Group**: Specifies the multicast group after NAT.

**Group Mask**: Specifies the mask value for the NAT rule. By default, it's 32.

**Pre-Translation Source**: Specifies the source IP address before NAT.

**Post-Translation Source**: Specifies the source IP address after NAT.

**Note**
The Post-Translation Source IP needs to be the secondary IP address on the loopback interface to make sure RPF check won't fail.

**Source Mask**: Specifies the source mask value for the NAT rule. By default, it's 32.

**Post-Translation Source Port**: Source Port is 0 by default. The value 0 means no translation.

**Post-Translation Destination Port**: Destination Port is 0 by default. The value 0 means no translation.

**Statis Oif**: This field is disabled for the **Ingress** mode. In the **Egress** mode, it populates the interfaces based on the Egress Interface Mappings defined.

**Step 7** Click **Save & Deploy** to save the NAT rule and deploy it.

### Deleting NAT Rule

**Procedure**

**Step 1** Choose **LAN** > **Fabrics**.

**Step 2** Double-click a fabric name.

The **Fabric Overview** window appears.

**Step 3** Click the **Multicast NAT** tab.

**Step 4** Click the **NAT Rules** tab.

**Step 5** Select the NAT mode that you need to delete and click **Actions** > **Delete NAT Rule** to delete a NAT rule.

If the NAT rule isn't deployed or failed, you can skip this step.

**Step 6** Click **Confirm** to delete the selected NAT rule.

# RTP/EDI Flow Monitor

**Note** This tab is only available on IPFM fabric when you have deployed IPFM on Nexus Dashboard Fabric Controller.

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **RTP/EDI Flow Monitor**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **RTP/EDI Flow Monitor**.

**Note** This section is applicable for both the IPFM and Generic Multicast modes in Nexus Dashboard Fabric Controller.

Cisco Nexus Dashboard Fabric Controller provides a view of all the active RTP and EDI streams. It also lists out active flows that have RTP and EDI drops and historical records for the same. For active IPFM flow, Nexus Dashboard Fabric Controller provides RTP and EDI topology to pinpoint the loss in network.

**Note** You need to enable telemetry in the switches to view RTP/EDI Flow Monitor. For more information, refer your respective platform documentation.

The description of the fields in these tabs are:

| Field | Description |
|---|---|
| VRF | Specifies the name of the VRF. |
| Switch | Specifies the name of the switch. |
| Interface | Specifies the interface from which the flows are detected. |
| Source IP | Specifies the source IP address of the flow. |
| Source Port | Specifies the source port of the flow. |
| Destination IP | Specifies the destination IP address of the flow. |
| Destination Port | Specifies the destination port of the flow. |
| Receiver IPs | Specifies the receiver IPs which are connected directly to the given switch. |
| Bit Rate | Specifies the bit rate of the flow, in bps, kbps, mbps, gbps, or tbp. |
| Packet Count | Specifies the number of packets in the flow. |
| Packet Loss | Specifies the number of lost packets. |
| Loss Start | Specifies the time at which the packet loss started. |
| Loss End | Specifies the time at which the packet loss stopped. |
| Start Time | Specifies the time at which the flow started. |
| Protocol | Specifies the protocol that is being used for the flow. |

You can click the **Telemetry Switch Sync Status** link to check whether the switches are in sync. The **Telemetry Sync Status** window displays the status of the switches in the **Sync Status** field and the last time that the sync occurred in the **Last Sync Time** field.

The RTP/EDI Flow monitor window has the following tabs:

- **Active Flows**
- **Packet Drop**
- **Drop History**

### Active Flows

The **Active Flows** tab displays the current active flows. You can also view these flows by navigating to **Flows > Flow Status**. You can click a switch link to view the end-to-end flow topology.

#### Flow Topology

The flow topology is displayed for the active flows that are displayed in the **Flow Status** window. For more information about multicast NAT visualization, see Flow Status.

From Cisco NDFC Release 12.1.2e, the flow topology for the active flows is displayed in the **Active Flows** tab.

Click a switch link to display the end-to-end flow topology.

The flow topology displays the direction of the flows. The arrows in the icon indicate the direction of the flow from the sender to the receiver. The IP addresses suffixed with **(S)** and **(R)** indicate the sender and receiver host respectively. If there are multiple receivers for a given flow, you can choose a receiver from the **Select Receiver** drop-down list.

The switches experiencing packet drops are circled in red.

Hover your cursor over a switch to display the following details:

- Name
- IP address
- Model
- Packet loss, if any

Click the **file** icon next to the links between the switches to view the interface counters errors for the interfaces connecting the two switches.

When you click the file icon, the **show interface <*interface name*> counters errors** command is run for the interface where the flow is participating between these switches, and the results are displayed in a pop-in.

### Packet Drop

The **Packet Drop** tab shows the packet drops for active flows.

### Drop History

When active RTP packet drop is not observed, records from the **Packet Drop** tab are moved to the **Drop History** tab. By default, the RTP drop history is maintained for 7 days. You can customize this setting by

entering the required value in the **IPFM history retention days** field in **Settings > Server Settings > IPFM** and saving it.

**Note** The **Drop History** tab displays only the last 100,000 records at the maximum.

# Global Config

**Note** This tab is only available on IPFM fabrics when you have deployed IPFM on Nexus Dashboard Fabric Controller. However, the IPFM fabric with generic multicast fabric technology is an exception (as the IPFM VRF created here is used for defining host/flow aliases for both IPFM and Generic Multicast Fabric).

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Global Config**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Global Config**.

Nexus Dashboard Fabric Controller allows two major operations.

- Monitor the network.
- Configure host and flow policies.

Nexus Dashboard Fabric Controller monitors the Flow Status, Discovered Host, Applied Host Policies, and other operations using Telemetry. For any operations triggered by the switch and received through telemetry (for example, Flow Established), Nexus Dashboard Fabric Controller periodically checks for new events and generate appropriate notification.

If `pmn.deploy-on-import-reload.enabled` server property is set to true during a switch reload, when Nexus Dashboard Fabric Controller receives switch coldStartSNMPtrap, it deploys Global Config, and Host and Flow policies that are showing 'Deployment Status=Successes' to the switch automatically. Deploy the switch telemetry and SNMP configuration can be deployed on demand by using Nexus Dashboard Fabric Controller packaged `pmn_telemetry_snmp` CLI template available in **Templates**.

Navigate to **Global Config** to set or modify Switch Global configuration and VRFs.

When you install Nexus Dashboard Fabric Controller with IPFM Deployment, you can deploy policies, the unicast bandwidth, Any Source Multicast (ASM) range, and VRFs using **Global Config**.

After you deploy the Nexus Dashboard Fabric Controller with IPFM, configure the bandwidth and ASM. The remaining percentage of the bandwidth is utilized by the multicast traffic. Nexus Dashboard Fabric Controller acts like a Master Controller, and deploy the bandwidth and ASM configurations to all the switches in the fabric.

As Cisco Nexus Dashboard Fabric Controller uses Telemetry to fetch data from the Fabric, the flow status and Kafka notifications may not reflect the current state in real time. It periodically checks new events and generates appropriate notification. For more information, refer to the *Kafka Notifications for Cisco Nexus Dashboard Fabric Controller, Release 12.0.1a*.

This section contains the following:

# Switch Global Config

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Global Config** > **Switch Global Config**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Global Config** > **Switch Global Config**.

Navigate to **Switch Global Config** to configure the global parameters.

**Note** A user with the network operator role in Nexus Dashboard Fabric Controller cannot save, deploy, undeploy, add or delete ASM, or edit the unicast bandwidth reservation percentage.

After you configure the Unicast Bandwidth Reservation and ASM range, you can perform the following operations to deploy these configurations to the switches.

After deploying the global configurations, configure the WAN for each switch in your network.

*Table 30: Switch Global Config Table Fields and Description*

| Field | Description |
|---|---|
| VRF | Specifies the name of the VRF. This VRF is used to associate IPFM Host/Flow policies as well as Host/Flow aliases for both IPFM and Generic Multicast fabrics. |
| Unicast Bandwidth Reservation % | Displays a numeric value that indicates the unicast bandwidth configuration percentage, and the status specifies if the bandwidth deployment was success, or failed or not deployed.<br>You can configure the server to allot a dedicated percentage of bandwidth to unicast traffic. The remaining percentage is automatically reserved for multicast traffic.<br>Click the numerical value link to view the details of the deployment history for the Unicast Bandwidth for the selected VRF and switch in the **Deployment History** pane. For more information, see Deployment History, on page 271.<br>Click the **Failed** or **Success** link to view the details of the deployment status for the Unicast Bandwidth for the selected VRF and switch in the **Deployment Status** pane. For more information, see Deployment Status, on page 271. |

| Field | Description |
|---|---|
| Reserve Bandwidth to Receiver Only | Bandwidth reservation status specifies if the bandwidth deployment was success, or failed or not deployed.<br>The Enabled status indicates that the ASM traffic is pushed to the spine only if there is a receiver. This feature is applicable for switches with the Cisco NX-OS Release 9.3(5) and later.<br>Click the **Enabled** link to view the details of the deployment history for the Reserve Bandwidth for the selected VRF and switch in the **Deployment History** pane. For more information, see Deployment History, on page 271.<br>Click the **Failed** link to view the details of the deployment status for the Reserve Bandwidth for the selected VRF and switch in the **Deployment Status** pane. For more information, see Deployment Status, on page 271. |
| ASM/MASK | Displays the number of Any Source Multicast (ASM) groups enabled for the selected VRF and the status indicates whether the ASM and Mask configuration was deployed successfully, or failed or not deployed.<br>The ASM is a PIM tree building mode that uses shared trees to discover new sources and receivers as well as source trees to form shortest paths from receivers to sources. ASM provides discovery of multicast sources.<br>The IP address and subnet mask in the ASM/MASK field define the multicast source.<br>The ASM range is configured by specifying the IP address and the subnet mask.<br>Click the numerical value link to view the details of the deployment history for the ASM/mask for the selected VRF and switch in the **Deployment History** pane. For more information, see Deployment History, on page 271.<br>Click the **Failed** link to view the details of the deployment status for the ASM/mask for the selected VRF and switch in the **Deployment Status** pane. For more information, see Deployment Status, on page 271. |

Click the table header to sort the entries in alphabetical order of that parameter.

The following table describes the action items, in the **Actions** drop-down list, that appears in the **Switch Global Config** window.

*Table 31: Switch Global Config Actions and Description*

| Action Item | Description |
|---|---|
| Edit NBM VRF Config | Allows you to edit the NBM VRF configuration.<br>To perform an edit, choose this option. The **Edit NBM VRF Config** window opens. Edit the required values and click **Deploy**. |
| Undeploy All | Undeploys ASM, unicast bandwidth, and reserved bandwidth configuration to all switches. |
| Undeploy Unicast BW | Undeploys only unicast bandwidth configuration. |
| Undeploy Reserve BW | Undeploys only the reserve bandwidth configuration. |
| Undeploy ASM/Mask | Undeploys only the ASM configuration. |
| Redo All Failed | Redeploys the selected failed configurations. |

### Deployment History

The following table describes the fields that appear on the Deployment History.

*Table 32: Deployment History Field and Description*

| Field | Description |
|---|---|
| Type | Specifies whether the type is Unicast Bandwidth Reservation %, Reserve Bandwidth to Receiver Only, or ASM/MASK. |
| VRF | Specifies the name of the VRF. |
| Switch Name | Specifies the switch name in the fabric on which the configuration was deployed. |
| Deployment Status | Displays the status of deployment. It shows if the deployment was **Success** or **Failed** along with the reason why the deployment failed. |
| Action | Specifies the action that is performed on the switch, such as **Create** or **Delete**. |
| Deployment Date/Time | Displays the date and time when the deployment was initialized. |

### Deployment Status

The following table describes the fields that appear on the Deployment Status.

*Table 33: Deployment Status Field and Description*

| Field | Description |
|---|---|
| Type | Specifies whether the type is Unicast Bandwidth Reservation %, Reserve Bandwidth to Receiver Only, or ASM/MASK. |
| VRF | Specifies the name of the VRF. |
| Switch Name | Specifies the switch name in the fabric on which the configuration was deployed. |
| IP Address | Specifies the IP address of the switch. |
| Deployment Status | Displays the status of deployment. It shows if the deployment was **Success** or **Failed** along with the reason why the VRF deployment failed. |
| Action | Specifies the action that is performed on the switch, for example, **Create**. |
| Deployment Date/Time | Displays the date and time when the deployment was initialized. |

# IPFM VRF

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Global Config** > **IPFM VRF**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Global Config** > **IPFM VRF**.

From Cisco NDFC Release 12.1.2e, **IPFM VRF** tab is included under the **Fabric Overview** window. Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **IPFM VRF**.

Use the **IPFM VRF** window to create, edit, delete, and redeploy IPFM VRFs. You can view the deployment status and history of each VRF.

From Cisco NDFC Release 12.1.2e, you can configure and monitor both NBM active and passive VRFs. In NBM passive mode, NDFC will be involved only in the monitoring of IPFM fabric and not configuration except in setting up VRF mode as NBM passive. Perform the following steps to change the NBM mode:

- Click **Actions** > **Create VRF**.
- On the **Create VRF** window, enter the name of the VRF. Choose **Active** or **Passive** and click **Save & Deploy**.



**Note** You cannot edit the existing VRF to change the NBM mode. You must delete and re-create VRF to change the NBM mode from active to passive or conversely.

If fabric is set to monitor mode, changing VRF is not applicable as this is fabric level configuration and not VRF configuration.

You are not allowed to create **IPFM VRF** when none of the switches are imported to NDFC. Import or add switch to the fabric to create IPFM VRF.

Discovery status is updated at regular interval by a background process. NBM configuration can be deployed even if the switch is in an unreachable state. After periodic discovery, the status of switches are updated appropriately.

***Table 34: IPFM VRF Table Fields and Description***

| Field | Description |
|---|---|
| Name | Specifies the name of the VRF. |
| Mode | Specifies the type of mode (Active or Passive) of the VRF. |
| Deployment Status | Specifies whether the VRF deployment is successful, failed, or the VRF is not deployed. For default VRFs, the deployment status is displayed as **Not Applicable**.<br>Click the **Failed** status to view more information about the Deployment Status, on page 271. |
| Deployment History | Specifies the deployment history of the VRF. For default VRFs, the deployment history is displayed as **Not Applicable**.<br>Click **View** in Deployment History to view more information about the Deployment History. |
| Description | Specifies the description of the VRF. |

Click the table header to sort the entries in alphabetical order of that parameter.

The following table describes the action items, in the **Actions** drop-down list that appears in the **IPFM VRF** horizontal tab on the **Global Config** tab in the **Fabric Overview** window.

*Table 35: IPFM VRF Actions and Description*

| Action Item | Description |
|---|---|
| Create VRF | Allows you to create a new VRF.<br>To create a VRF, choose **Create VRF** from the **Action** drop-down list of the **IPFM VRF** horizontal tab on the **Global Config** tab in the **Fabric Overview** window. In the **Create VRF** window, enter the VRF name and description, choose **Active** or **Passive** mode and click **Save & Deploy** to retain the changes and deploy or click **Cancel** to discard the changes.<br>**Note**<br>When you create an active nondefault VRF, although the default host and flow policies are automatically created for that VRF, you must manually deploy the policies to the switches in the fabric. When VRF is set to passive, then flow policies are not created. For more information about deploying the policies manually, see Host Policies, on page 225 and Flow Policies. |
| Edit VRF | Allows you to edit a selected VRF.<br>To edit a VRF, select the check box next to the VRF that you want to edit and choose **Edit VRF**. In the **Edit VRF** window, you can edit only the description and click **Save** to retain the changes or click **Cancel** to discard the changes. |
| Delete VRF | Allows you to delete one or more VRFs, which deletes the data from the database and cancels the deployment on the switch.<br>To delete a VRF, select the check box next to the VRF that you want to delete and choose **Delete VRF**. You can select multiple VRF entries and delete them at the same instance. |
| Redeploy | Allows you to select and redeploy the VRFs with failed status.<br>To redeploy a VRF to the switch, select the check box next to the VRF that you want to deploy again and choose **Redeploy**. You can select multiple VRF entries and redeploy them at the same instance. |

### Deployment History

The following table describes the fields that appear in the **Deployment History** pane.

*Table 36: Deployment History Field and Description*

| Field | Description |
|---|---|
| Type | Specifies the type of VRF. |
| VRF | Specifies the name of the VRF. |
| Switch Name | Specifies the switch on which the VRF is deployed. |

| Field | Description |
|---|---|
| Deployment Status | Displays the status of the deployment. It shows if the deployment was a **Success**, **Failed** along with the reason why the VRF deployment failed, or **Not Applicable**. |
| Action | Specifies the action that is performed on the switch, such as **Create** or **Delete**. |
| Deployment Date/Time | Displays the date and time when the deployment was initialized. |

### Deployment Status

The following table describes the fields that appear in the **Deployment Status** pane.

*Table 37: Deployment Status Field and Description*

| Field | Description |
|---|---|
| Type | Specifies the type of VRF. |
| VRF | Specifies the name of the VRF. |
| Switch Name | Specifies the switch on which the VRF is deployed. |
| IP Address | Specifies the IP address of the switch. |
| Deployment Status | Displays the status of the deployment. It shows if the deployment was a **Success** or **Failed** along with the reason why the deployment failed. |
| Action | Specifies the action that is performed on the switch, for example, **Create**. |
| Deployment Date/Time | Displays the date and time when the deployment was initialized. |

# VRF (Generic Multicast)

**Note** This tab is only available on IPFM fabric when you have deployed IPFM on Nexus Dashboard Fabric Controller and when the fabric technology is generic multicast.

**UI Navigation**

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **VRF**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **VRF**.

Use the **VRF** window to create, edit, and delete VRFs.

*Table 38: VRF Table Fields and Description*

| Field | Description |
|---|---|
| Name | Specifies the name of the VRF. |
| Deployment Status | For generic multicast VRFs, the deployment status is displayed as **Not Applicable**. |
| Deployment History | For generic multicast VRFs, the deployment status is displayed as **Not Applicable**. |
| Description | Specifies the description of the VRF. |

Click the table header to sort the entries in alphabetical order of that parameter.

The following table describes the action items, in the **Actions** drop-down list, that appears in the **VRF** window.

*Table 39: VRF Actions and Description*

| Action Item | Description |
|---|---|
| Create VRF | Allows you to create a new VRF.<br>To create a VRF, choose **Create VRF** from the **Action** drop-down list on the **VRF** tab in the **Fabric Overview** window. In the **Add VRF** window, enter the VRF name and description, and click **Save** to retain the changes or click **Cancel** to discard the changes. |
| Edit VRF | Allows you to edit a selected VRF.<br>To edit a VRF, select the check box next to the VRF that you want to edit and choose **Edit VRF**. In the **Edit VRF** window, you can edit only the description and click **Save** to retain the changes or click **Cancel** to discard the changes. |
| Delete VRF | Allows you to delete a selected VRF.<br>To delete a VRF, select the check box next to the VRF that you want to delete and choose **Delete VRF**. You can select multiple VRF entries and delete them at the same instance. |

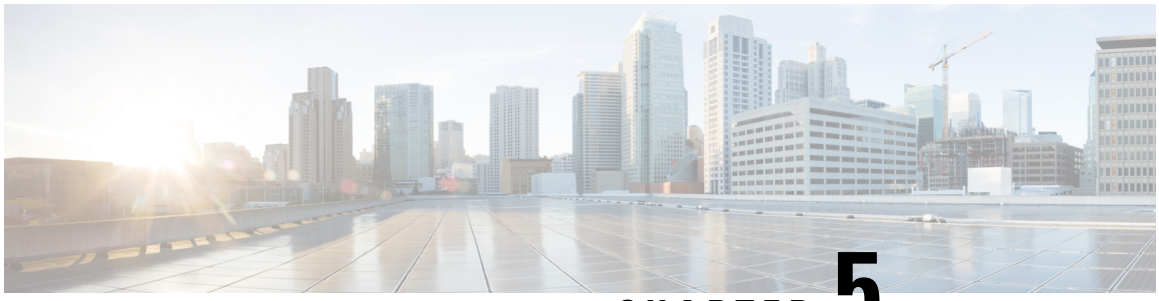
# Virtual Infrastructure

### Viewing OpenStack VMs

The following table describes the fields and description on the window.

| Field | Description |
|---|---|
| VM Name | Specifies the name of the Kubernetes pod. |
| Compute Name | Displays the IP address of the Kubernetes pod. |

| Field | Description |
|---|---|
| Fabric Name | Specifies the phase (state) of the pod. |
| IP Address | Specifies the reason. |
| MAC Address | Specifies the applications of the pod. |
| Physical NIC | Specifies the namespace of the pod. |
| Port Channel | Specifies the node name of the pod. |
| Switch Interface | Specifies the switch interface connected to pod. |
| Switch Name | Specifies the name of the switch. |
| Switch IP | Specifies the IP address of the switch. |
| VLAN | Specifies the VLAN. |
| Locked | Specifies the whether the cluster is in locked state. |
| Power State | Specifies whether the openstack cluster power state. |
| Network State | Specifies whether the openstack cluster network state. |
| State | Specifies the state of openstack cluster. |



CHAPTER 5

# Switches


- Switches, on page 279
- Switch Overview, on page 303


## Switches

The following table describes the fields that appear on **Switches** window.

| Field | Description |
|---|---|
| Switch | Specifies name of the switch. |
| IP Address | Specifies IP address of the switch. |
| Role | Specifies role assigned on the switch. |
| Serial Number | Specifies the serial number of the switch. |
| Fabric Name | Specifies the associated fabric name for the switch. |
| Config Status | Specifies the configuration status.<br>Status will be either In-Sync or Out-of-sync. |
| Oper Status | Specifies the configuration status.<br>Status will be either In-Sync or Out-of-sync. |
| Discovery Status | Specifies the discovery status of the switch. |
| Model | Specifies the switch model. |
| vPC Role | Specifies the vPC role of the switch. |
| vPC Peer | Specifies the vPC peer of the switch. |

## Adding Switches to a Fabric

UI Path: **LAN** > **Switches** > **Actions** > **Add Switches**

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



**Note** Cisco Nexus Dashboard has 2 logical interfaces per node, namely, Management interface (bond1br) and Fabric (also known as data) interface (bond0br). For Cisco Nexus Dashboard Fabric Controller, Nexus Dashboard Management and Fabric interfaces must be in different IP subnets. By default, the route for Nexus Dashboard services is through the fabric interface. An operator must add static routes on Nexus Dashboard Management Network to connect with switches that must be reached over management interface (bond1br). This ensures that a route for the pods uses management interface as the exit interface.



**Note** Make sure that the switch user role for discovery or add switches or LAN credentials for NDFC must have the network-admin role.

To add switches to the existing fabric, perform below procedures:

1. From Nexus Dashboard Fabric Controller Web UI, choose **LAN** > **Switches**.
2. On Switches tab, Choose **Actions** > **Add Switches**.

   The **Add Switches** window appears.

   Similarly, you can add switches on Topology window. On topology window, choose a fabric, right-click on the fabric and click **Add Switches**.
3. On add switches window, click **Choose Fabric**, click appropriate fabric, and then click **Select**.

   The **Add Switches** window has a default discover tab and other tabs appears based on the fabric selected.

Also, you can pre-provision switches and interfaces. For more information, see pre-provision device and pre-provisioning ethernet interface.



**Note** NDFC supports switch discovery only for default system-name(serial number).



**Note** When Nexus Dashboard Fabric Controller discovers a switch with the hostname containing the period character (.), it is treated as a domain-name and truncated. Only the text before the period character (.) is considered as a hostname. For example:

- If hostname is **leaf.it.vxlan.bgp.org1-XYZ**, Nexus Dashboard Fabric Controller shows only **leaf**
- If hostname is **leaf-itvxlan.bgp.org1-XYZ**, Nexus Dashboard Fabric Controller shows only **leafit-vxlan**



**Note** Ensure that the Switch name or the Host name is unique within the Fabric.

## 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 Nexus Dashboard Fabric Controller, the DHCP request from the device, will be forwarded to the Nexus Dashboard Fabric Controller. For easy day-0 device bring-up, the bootstrap options should be enabled on the **Fabric Settings** as mentioned earlier.
3. With bootstrap enabled for the fabric, the DHCP request coming from the device will be serviced by the Nexus Dashboard Fabric Controller. The temporary IP address allocated to the device by the Nexus Dashboard Fabric Controller will be employed to learn basic information about the switch including the device model, device NX-OS version, etc.
4. In the Nexus Dashboard Fabric Controller UI, choose **Switch** > **Actions** > **Add Switches**.

   The **Add Switches** window appears with default tabs.
5. Choose **Bootstrap(POAP)** radio button.

   As mentioned earlier, Nexus Dashboard Fabric Controller 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.

Note: You can change serial number only for Nexus 9000 Series switches.

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

Add the IPv4 address in the **IP Address** field.

You can provision devices in advance. To pre-provision devices, refer to Pre-provisioning device section.

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.

   You can specify a new user. Choose radio button **Specify a new user** enter **Username**, **Password** and choose **Authentication Protocol** from drop-down list.



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

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

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

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

If the discovery credentials are not provided, Nexus Dashboard Fabric Controller uses the admin user and password to discover switches.

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

Nexus Dashboard Fabric Controller 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 Deploy Config 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 Out-of-Sync, 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.
- When discovering devices using SNMP, if you have configured to use an AAA server for authentication, the command **sync-snmp-password** <*password*> <*username*> is run on the switch through NDFC to generate a cached user. The authentication uses MD5, by default. You must specify the SNMPv3 authentication and privacy protocol attributes in the switch AV-pair as follows: `snmpv3:auth=SHA priv=AES-128`

**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 Nexus Dashboard Fabric Controller. 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 Nexus Dashboard Fabric Controller 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 **LAN** > **Switches**. Select a switch, a slide-in pane appears, click **Launch** icon. On **Switches Overview** tab, click **Interface** tab 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 **Deploy Config** operation. At the end of the operation, an error prompts you to configure the **shutdown** or **no shutdown** command on the nve interface. A sample error screenshot when you enable a vPC setup.

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

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

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

**Note**

- Changing of the switch role is allowed only before executing **Deploy Config**.
- 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 section.

- **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 history.
- **History** - View per switch deployment and policy change history.

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

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

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

| PTI Operations | Generated Config Before | Generated Config After |
|---|---|---|
| Add | Empty | Contains the config |
| Update | Contains config before changes | Contains config after changes |
| Mark-Delete | Contains the config to be removed. | Contains the config to be removed with color 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 Nexus Dashboard Fabric Controller, the underlay configuration provisioned on those switches, and the configurations between Nexus Dashboard Fabric Controller and the switches are synced. The remaining tasks are:

- Provision interface configurations such as vPCs, loopback interface, and subinterface configurations.
- Create networks and deploy them on the switches.

## Discovering Existing Switches

To discover existing switches in Cisco Nexus Dashboard Fabric Controller Web UI, perform the following procedure:

### Procedure

**Step 1** After you click **Add Switches**, click **Discover Switches** 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.

**Step 2** The IP address (Seed IP), username, and password (**Username** and **Password** fields) of the switch are provided as the input by a user. The **Preserve Config** check box is chosen 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** check box is not selected.

**Note**
BGP Fabric does not support brownfield import of a device into the fabric.

**Step 3** Click **Discover Switches**.

The **Add Switches** window appears. 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 **Add Switches** result.

**Step 4** If the Cisco Nexus Dashboard Fabric Controller was able to perform a successful shallow discovery to a switch, the status column shows as **Manageable**. Choose the check box next to the appropriate switch(es) and click **Add Switches**.

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.

Cisco Nexus Dashboard Fabric Controller 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.

**Step 5** 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.

**Step 6** After discovering the devices, assign an appropriate role to each device. For more information on roles, refer Assign Roles.

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 Nexus Dashboard Fabric Controller, 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.

**Step 7** Click **Save**.

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.

**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 Cisco Nexus Dashboard Fabric Controller to the fabric are accurate or to detect any deviations (such as out-of-band changes), Nexus Dashboard Fabric Controller's Configuration Compliance engine reports and provides necessary remediation configurations.

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

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

The Re-sync button is displayed for each switch in the Re-sync column. Use this option to resynchronize Nexus Dashboard Fabric Controller state when there is a large scale out-of-band change, or if configuration changes do not register in the Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller.

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

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

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

Multi-line banner motd configuration can be configured in Cisco Nexus Dashboard Fabric Controller 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 **Deploy Config** 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.

**Step 8** Close the **screen**.

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 Nexus Dashboard Fabric Controller 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 **Deploy Config** 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.

An example of the **Deploy Config** option usage is for switch-level freeform configurations. Refer for details.

## Adding Switches Using Bootstrap Mechanism

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.

Starting from Nexus Dashboard Fabric Controller Release 12.0.1a, POAP access user validated key exchange and password-less ssh to limit configuration file access to the specific switch for a finite time. Therefore, you must accept a new key via **Add Switch > Bootstrap** whenever a device attempts POAP.

If there is IP reachability between the device and the Nexus Dashboard Fabric Controller, the DHCP request from the device, will be forwarded to the Nexus Dashboard Fabric Controller. For easy day-0 device bring-up, the bootstrap options should be enabled in the Fabric Settings.

With bootstrap enabled for the fabric, the DHCP request coming from the device will be serviced by the Nexus Dashboard Fabric Controller. The temporary IP address allocated to the device by the Nexus Dashboard Fabric Controller will be employed to learn basic information about the switch including the device model, device NX-OS version, etc.

1. Choose **LAN > Switches > Add Switches**.
2. Choose **Bootstrap(POAP)** radio button.
3. Click **Actions** and add Switches.

   You can add switches one at a time using the **Add** option or add multiple switches at the same time using the **Import** option.

   If you use the **Add** option, ensure you enter all the required details.

   Note: It might take some time for the switches to appear.
4. Choose a required switch.
5. Click **Edit**.

   The **Edit bootstrap switch** dialog appears.
6. Enter the required details.
7. Click **Save**.
8. Choose the switch.
9. Enter the admin password in the **Admin password** field.
10. Click **Import Selected Switches**.

## Return Material Authorization (RMA)

This section describes how to replace a physical switch in a Fabric when using Cisco Nexus Dashboard Fabric Controller 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 Recalculate config 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, Nexus Dashboard Fabric Controller pushes the **system mode maintenance** command to the switches when the Nexus Dashboard Fabric Controller RMA procedure is initiated. This command applies the configuration that is present in the default maintenance mode profile to the switches. For more information on performing Graceful Insertion and Removal (GIR) on the Cisco Nexus 7000 Series switches, refer Configuring GIR.

## POAP RMA Flow

To provision RMA, follow below procedure:

### Procedure

**Step 1** Navigate to the Fabric overview.

**Step 2** Move the device into maintenance mode. To move a device into maintenance mode, choose the device, click **Actions > More > Change Mode**. From the Mode drop-down list, select **Maintenance**.

**Step 3** 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 4** Power on the switch and go through the POAP cycle.

**Step 5** Initiate the RMA flow. Choose the device, click **Actions > More > Provision RMA**.

**Step 6** Set the admin password.

(Optional) You can set a AAA user and password for discovery.

**Step 7** Select the replacement device.

**Step 8** Click **Provision RMA**.

## Manual RMA Flow

Use this flow when Bootstrap is not possible (or not desired).

To provision manual RMA, follow below procedure:

### Procedure

**Step 1** (Optional) Place the device in maintenance mode.

**Step 2** Physically replace the device in the network.

**Step 3** Log in through Console and set the Management IP and credentials.

**Step 4** If you are using AAA, configure AAA commands on the switch.

Ensure you update LAN and discovery credentials in NDFC for the newly configured AAA user, if configured.

**Step 5** The Cisco Nexus Dashboard Fabric Controller rediscovers the new device or you can manually choose **Discovery** > **Rediscover**.

**Step 6** Deploy the expected configuration using **Actions** > **Deploy**.

**Step 7** Depending on the configuration, if breakout ports or FEX ports are in use, you have to deploy again to completely restore the configuration.

**Step 8** After a successful deployment, and the device is "In-Sync," you must move the device back to Normal Mode.

### 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 Nexus Dashboard Fabric Controller `switch_snmp_user` policy for the switch with the new SNMP MD5 key from the switch.

# Pre-provisioning Support

Cisco NDFC 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 NDFC prior to the device connectivity to the Network. Pre-provisioning is supported for Cisco NX-OS devices in Data Center VXLAN EVPN, External Connectivity Network, and Classic LAN fabrics.

## Pre-provisioning a Device

You can provision devices before adding them to fabrics. However, ensure that you enter DHCP details in the Bootstrap tab in the fabric settings.

The pre-provisioned devices support the following configurations in Nexus Dashboard Fabric Controller:

- 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 Nexus Dashboard Fabric Controller:

- 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 interface breakout CLI in the **Data** key of the pre-provision payload must contain the exact format as is on the 'show running-configuration' output from the switch.

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 Nexus Dashboard Fabric Controller 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"}

1. Choose **LAN** > **Switches**> **Add Switches**.
2. Choose **Pre-provision** radio button.
3. Click **Actions** and add switches.

   You can add switches one at a time using the **Add** option or add multiple switches at the same time using the **Import** option.

   If you use the **Add** option, ensure you enter all the required details.
4. Choose a switch.
5. Enter the admin password in the **Admin password** field.
6. Click **Pre-provision**.

   The pre-provisioned switch is added.

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

   For more information, see Return Material Authorization (RMA).

   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.

## Pre-provisioning an Ethernet Interface

You can pre-provision Ethernet interfaces in the **LAN Interfaces** 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 NDFC.

**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 pre-provisioned device in your fabric. For information, see Pre-provisioning a Device, on page 289.

### Procedure

**Step 1** Double-click on the fabric containing the pre-provisioned device from the **LAN Fabrics** window.

The **Fabric Overview** window appears.

**Step 2** On the **Interfaces** tab, click **Actions** > **Create Interface**.

The **Create Interface** window appears.

**Step 3** Enter all the required details in the **Create Interface** window.

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

**Select a device**: Select the pre-provisioned device.

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

**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 323.

**Step 4** Click **Save**.

**Step 5** 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 more information, refer Pre-provisioning a Device, on page 289.

Two Cisco Nexus 9000 Series devices that are pre-provisioned and added to an existing Fabric. Choose **Add Switches** from the **Actions** drop-down list. On the Inventory Management screen, click **PowerOn Auto Provisioning (POAP)**.

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

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

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

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

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

For instructions, see Pre-provisioning an Ethernet Interface, on page 291.

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

In the **Links** tab, click on **Actions** > **Create**.

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.

After the links are created, they are listed in the **Links** tab under **Fabric Overview** window.

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

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

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

After completion, the devices will be correctly paired and the vPC pairing intent will be generated for the devices and the policies are generated

**Note**
Because the devices are not yet operational, Configuration 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 more information, see Pre-provisioning an Ethernet Interface, on page 291.

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

**Step 2** Create a vPC host truck interface .

**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**.

## Attaching Overlays to Pre-provisioned Devices

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

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 the **Fabric Overview** window, click the **Policies** tab and choose **Actions** > **Edit Policy**. You can view the entire intent generated for the pre-provisioned device, including the overlay network/VRF attachment information.

# Previewing Switches

**Nexus Dashboard Fabric Controller UI Navigation**

- Choose **LAN** > **Switches**.
- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric Summary** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Switches**.

After adding the switches, you can preview the switches with pending configurations, the side-by-side comparison of running configurations, and the expected configurations for the switches. You can select multiple switches and preview them at the same instance. The **Preview** window displays the pending configurations for the successful deployment of a switch.

To preview the switches and resync the ones with pending configurations, perform the following steps:

**Procedure**

**Step 1** In the **Switches** window, use the check boxes next to the switches to select the switches that you want to preview. From the **Actions** drop-down list, choose **Preview**.

The **Preview Config** window appears. This window displays the switch configuration information such as the switch name; its ip address, role, serial number; the fabric status-whether it is in sync, out of sync, or not available; the pending configuration; the status description; and the progress.

**Step 2** To only preview the configuration, view the displayed information and click **Close**.

**Step 3** To resynchronize the switches with pending configuration, click **Resync**. The progress bar displays the progress of the resynchronization. Click **Close** to close the **Preview Config** window.

**Step 4** To view the pending configurations and side-by-side comparison, click the respective link in the **Pending Config** column.

Alternatively, on the **Fabric Overview Actions** drop-down list, select **Recalculate Config**. The **Deploy Configuration** window appears. It displays the configuration status on the switches. You can also view the pending configurations by clicking the respective link in the **Pending Config** column.

The **Pending Config** window appears. The **Pending Config** tab on this window displays the pending configurations on the switch. The **Side-by-Side Comparison** tab displays the running configuration and expected configuration side-by-side.

Close the **Pending Config** window.

# Deploy Configuration

This deploy option is a local operation for a switch, that is, the expected configuration or intent for a switch is evaluated against it’s current running configuration, and a config compliance check is performed for the

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

1. Choose required switch, choose **Actions** > **Deploy** to deploy configuration on a switch.

   The **Deploy Configuration** window appears.

2. Click **Resync** to synchronize configuration.
3. Click **Deploy**.

   The Status column displays FAILED or SUCCESS state. For a FAILED status, investigate the reason for failure to address the issue.

4. Click **Close** to navigate to switch window.

# Discovery

This chapter contains below sections:

## Update Credentials

Use update discovery credentials for updating discovering switches.

**Procedure**

**Step 1** Choose required switch, choose **Actions** > **Discovery** > **Update Credentials**.

The **Update Discovery Credentials** window appears.

**Step 2** In the **Update Discovery Credentials** window, enter the discovery credentials such as discovery username and password.

**Step 3** Click **Update** to save the discovery credentials.

If the discovery credentials are not provided, Nexus Dashboard Fabric Controller uses the admin user and password to discover switches.

## Rediscover

You can rediscover switch and check the status of it.

To rediscover the switch:

- Choose required switch, choose **Actions** > **Discovery** > **Rediscover** to rediscover switches.

  The **Discovery Status** column shows the status as **Rediscovering** and after discovering it displays the status.

# Guidelines and Limitations for Changing Discovery IP Address

From Cisco Nexus Dashboard Fabric Controller Release 12.0.1a, you can change the Discovery IP address of a device that is existing in a fabric.

### Guidelines and Limitations

The following are the guidelines and limitations for changing discovery IP address.

- Changing discovery IP address is supported for NX-OS switches and devices that are discovered over their management interface.
- Changing discovery IP address is supported for templates such as:
  - Easy_Fabric
  - Easy_Fabric_eBGP
  - External
  - LAN_Classic
  - LAN_Monitor
- Changing discovery IP address is supported in both managed and monitored modes.
- Only users with the **network-admin** role can change the discovery IP address on Cisco Fabric Controller UI.
- The discovery IP address must not be used on other devices, and it must be reachable when the change is done.
- While changing the discovery IP address for a device in a managed fabric, switches are placed in migration mode.
- When you change the IP address of a switch that is linked to vPC Peer, corresponding changes such as vPC peer, domain configuration will be updated accordingly.
- Fabric configuration restores the original IP address, it reports out of sync post restore and the configuration intent for the device must be updated manually to get the in-sync status.
- Fabric controllers restore that had the original device discovery IP reports the switch as Unreachable post restore. The discovery IP address change procedure must be repeated after the restore.
- Device Alarms associated with the original discovery IP address will be purged after the change of IP address.

## Changing Discovery IP Address

### Before you begin

You must make the management IP address and route related changes on the device and ensure that the reachability of the device from Nexus Dashboard Fabric Controller.

To change the discovery IP address from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN** > **Fabrics**.

**Step 2** Click on fabric names to view the required switch.

The **Fabric summary** slide-in pane appears.

**Step 3** Click **Launch** icon to view **Fabric Overview** window.

**Step 4** On the **Switches** tab, click **Refresh** icon adjacent to the **Action** button on the main window.

Switch with a changed IP address will be in **Unreachable** state in **Discovery Status** column.

**Step 5** Click the check box next to the **Switch** column and select the switch.

**Note**
You can change the IP address for individual switch and not for multiple switches.

**Step 6** Choose **Actions** > **Change Discovery IP** on the switches tab area.

The **Change Discovery IP** window appears.

Similarly, you can navigate from **LAN** > **Switches** tab. Choose a required switch, click **Actions** > **Discovery** > **Change Discovery IP**.

**Step 7** Enter the appropriate IP address in the **New IP Address** text field and click **OK**.

a) The new IP address must be reachable from Nexus Dashboard Fabric Controller to update successfully.
b) Repeat the above procedures for the devices where the discovery IP address must be changed before proceeding with further steps.
c) If the fabric is in managed mode, the device mode will be updated to migration mode.

**Step 8** From the fabric **Actions** drop-down list, click **Recalculate Config** to initiate the process of updating Nexus Dashboard Fabric Controller configuration intent for the devices. Similarly, you can recalculate configuration on topology window. Choose **Topology**, tab right-click on the switch, click **Recalculate Config**.

The Nexus Dashboard Fabric Controller configuration intent for the device management related configuration will be updated and the device mode status for the switch is changed to normal mode. The switch configuration status is displayed as **In-Sync**.

**Note**
The PM records associated with the old switch IP address will be purged and new record collections take an hour to initiate after the changes.

# Assigning Switch Roles

You can assign roles to switches on Nexus Dashboard Fabric Controller.

1. Choose required switch, choose **Actions** > **Set Role**.
2. The **Select Role** window appears. You can choose appropriate role and click **Select**.

   A confirmation window appears.

**Note** You must rediscover the switch to view new role assignation in **Role Status** column.

The following roles are supported in Nexus Dashboard Fabric Controller:

- Spine
- Leaf
- Border
- Border Spine
- Border gateway
- Border gateway spine
- Super spine
- Border super spine
- Border gateway super spine
- Access
- Aggregation
- Edge router
- Core router
- TOR

# Creating a vPC Setup

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

### Procedure

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

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

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

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

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

**vPC Domain tab**: Enter the vPC domain details.

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

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

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

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

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

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

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

**Step 3** Click **Save**.

The **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.

After creating a vPC pair, you can view vPC details in **vPC Overview** window.

## 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 **Deploy Config**.

**Step 4** (Optional) Click the value under the **Recalculate 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 **Deploy Config**: 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.

# Performing Actions on Switches

### Change Mode

To change mode for the switch, perform the following steps:

1. Choose check box for required switch, choose **Actions** > **More** > **Change Mode**.

   The **Change Mode** window appears.

2. Choose required **Normal** or **Maintenance** from drop-down list.
3. Click **Save and Deploy Now** to change mode or click **Save and Deploy Later** to change mode later.

### Provision RMA

To change mode for the switch, perform the following steps:

1. Choose check box for required switch, choose **Actions** > **More** > **Provision RMA**.

   The **Provision RMA** window appears.

2. The Provision RMA UI will show the replacement device 5-10 minutes after it is powered on.

### Copy Run Start

To copy the existing switch configuration to start configuration, perform the following steps:

1. Choose check box for required switch, choose **Actions** > **More** > **Copy Run Start**.

   The **Copy Running Config to Startup Config** screen appears. In the Progress column shows the process in progress and status description shows **Deployment in progress**.

2. A confirmation window appears, click **OK**.

The status description column displays **Deployment completed** and progress column in green.

3. Click **Close** to close this window.

### Reload

To reload required switch, choose **Actions** > **More** > **Reload**.

A confirmation window appears, click **Confirm**.

### Restore Switch

You can restore a Cisco Nexus switch in external fabrics and LAN classic fabrics from the Cisco Nexus Dashboard Fabric Controller 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.

1. Choose **Actions** > **More** > **Reload**.

   The **Restore Switch** window appears and you are in the **Select a Backup** tab. Refer to Backup Fabric for more information.

2. The **Select a Backup** tab displays the fabric backup details. It includes the following information:

   - Backup Date - Specifies the backup date and time.
   - Backup Version - Specifies the version number of backup.
   - Backup Tag - Specifies the name of backup.
   - NDFC Version - Specifies the NDFC version details.
   - Backup Type - Specifies the type of backup, either manual or automatic.

     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.

     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 Nexus Dashboard Fabric Controller archives only up to 10 golden backups.

3. Choose radio button for required backup to mark as golden, choose **Actions > Mark as golden**, a confirmation window appears, click **Confirm**.

4. Choose radio button for backup to delete from golden, choose **Actions > Remove as golden**, a confirmation window appears, click **Confirm**.

**Note** Most of this information is at the fabric level, and may or may not directly impact the proceedings of the switch-level restore.

5. Click **Next** to move to the **Restore Preview** step.
6. You can view information about the switch name, switch serial, IP address, status, restore supported, delta configuration and the VRF details.
7. (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.

8. Click **Restore Intent** to proceed to the **Restore Status** step in restoring.

   The restore status and description appears for the switch.
9. Click **Finish** 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.

### Show Commands

The following procedure view the commands in Nexus Dashboard Fabric Controller:

1. Choose **Actions** > **More** > **Show commands**.

   The **Switch Show Commands** window appears.
2. Choose required commands from drop-down list and enter required information in text field.
3. Click **Execute** to view the CLI output and to clear the output, click **Clear Output**.

### Exec Commands

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 Nexus Dashboard Fabric Controller:

1. Choose **Actions** > **More** > **Exec commands**.

   The **Switch Show Commands** window appears.

2. From the **Template** drop-down list, choose **exec_freeform** or **exec_elam_capture**.
3. Enter the commands in the **Freeform CLI** for **exec_freeform** and required IP addresses.
4. Click **Deploy** to run the EXEC commands.
5. In the **CLI Execution Status** window, you can check the status of the deployment. Click **Detailed Status** under the **Command** column to view details.
6. In the **Command Execution Details** window, click the info under the **CLI Response** column to view the output or response.

### Delete Switches

You can delete one or more existing switches.

Choose **Actions** > **More** > **Delete switche(s)**. A confirmation window appears, click **Confirm**

# Switch Overview

You can perform below operations, from **Actions** icon on Switch Overview window:

- Preview
- Deploy
- Discovery
- Set Roles
- vPC Pairing
- More

## Viewing Switch Overview

You can view information about switch along with the switch summary on **Switch Overview** tab. Navigate **LAN** > **Switches**, click on required switch. A slide-in pane appears. Click **Launch** icon to view the **Switch Overview** window.

You can search required details in **Filter by attributes** for multiple queires. For an example with Switch Names and Switch IP show in table

| Switch | IP Address |
|---|---|
| leaf1 | 10.10.10.1 |
| leaf2 | 10.10.10.2 |
| spine1 | 10.10.11.1 |
| spine2 | 10.10.11.2 |

When you search for Query String: **Switch Name == leaf1 Switch Name == leaf2**, It displays leaf1 and leaf2 switch name and IP.

When you search for Query String: **Switch Name != leaf1 Switch Name != leaf2**, It displays all switch names and its details.

| Field | Description |
|---|---|
| Switch Info | Specifies the switch information such as switch name, IP address, switch model and other details. |
| Alarms | Specifies the alarms configured on the selected switch |
| Performance | Specifies the CPU utilization and memory utilization for the switch. |
| Interfaces | Specifies the interface details. |
| Modules/FEX | Specifies the modules and FEX information. |
| Reports | Specifies the reports. |

# Hardware

This tab contains below sections:

## Modules

To view the inventory information for modules from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Procedure

**Step 1** Choose **LAN > Switch > Switch Overview > Hardware > Modules**.

The **Modules** tab is displayed with a list of all the switches and its details for a selected Scope.

You can view required information in table, enter details in **Filter by Attributes**.

**Step 2** You can view the following information.

- **Name** displays the module name.
- **Model** displays the model name.
- **Serial Number** column displays the serial number.
- **Type** column displays the type of the module.
- **Oper. Status** column displays the operation status of the module.
- **Slot** column displays the slot number.
- **HW Revision** column displays the hardware version of the module.

- **Software Revision** column displays the software version of the module.
- **Asset ID** column displays the asset id of the module.

## Viewing Bootflash

You can view the following information on Bootflash tab.

- **Primary Bootflash Summary** card displays the total, used and available space.
- **Secondary Bootflash Summary** card displays the total, used and available space.
- **Directory Listing** area displays check box for **Primary Bootflash** and **Secondary Bootflash**.

  This area shows the filename, size, and last modified date for all the files and directories on the switch bootflash. Choose **Actions** > **Delete** to delete files to increase the available space on the switch.

# 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 Nexus Dashboard Fabric Controller.

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 color 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**.

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.

Starting from Cisco NDFC Release 12.1.2e, parameters MTU, SPEED, Source Interface Description, Destination Interface Description, Source Interface Freeform Config, and Destination Interface Freeform Config are added to the existing **int_pre_provision_intra_fabric_link** template. These parameters are preserved on subsequent **Recalculate & Deploy** after the device has completed bootstrap and POAP.

The following table describes the fields that appear on **Links** tab.

| Field | Description |
|---|---|
| Fabric Name | Specifies the name of the Fabric. |
| Name | Specifies the name of the link.<br>The list of previously created links is displayed. The list contains intra-fabric links, which are between switches within a fabric, and inter-fabric links, which are between border switches in this fabric and switches in other fabrics. |

| Field | Description |
|---|---|
| Policy | Specifies the link policy. |
| Info | Provides more information about the link. |
| Admin State | Displays the administrative state of the link. |
| Oper State | Displays the operational state of the link. |

The following table describes the action items, in the Actions menu drop-down list, that appear on **Fabric Overview** > **Links** > **Links**.

| Action Item | Description |
|---|---|
| Create | Allows you to create the following links:<br>• Creating Inter-Fabric Links, on page 191<br>• Creating Intra-Fabric Links, on page 190 |
| Edit | Allows you to edit the selected fabric. |
| Delete | Allows you to delete the selected fabric. |
| Import | 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.<br>**Note**<br>• You cannot update existing links.<br>• The Import Links icon is disabled for external fabric. |
| Export | Choose the link and select **Export** to export the links in a CSV file.<br>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. |

# PTP (Monitoring)



**Note** PTP Monitoring can be installed as an application, and this application works only in IPFM mode.

**UI Navigation**

- Choose **LAN** > **Switches**. Click on a switch to open the **Switch** slide-in pane. Click the **Launch** icon. Choose **Switch Overview** > **PTP**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Switches**. Double-click a switch to open **Switch Overview** > **PTP**.
- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric Summary** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Switches**. Click on a switch to open the **Switch** slide-in pane and then click the **Launch** icon. Alternatively, you can double-click a switch to open **Switch Overview**. Choose **Switch Overview** > **PTP**.

This section explains the preview functionality of the Precision Time Protocol (PTP) monitoring. PTP is a time synchronization protocol for nodes that are distributed across a network. On a local area network, it achieves clock accuracy in the sub-nanosecond range, making it suitable for measurement and control systems.

On the **PTP** tab in the **Switch Overview** window, you can view PTP-related information based on the selected switch. You can click the **Telemetry Switch Sync Status** link to check whether the switches are in sync. The **Sync Status** column displays the status of the switches.

The following tabs are displayed in this window:

- **Correction & Mean Path Delay**
- **Clock & Port Status**

### Correction and Mean Path Delay

The **Correction & Mean Path Delay** tab displays a graph showing the PTP operational statistics: mean path delay, correction, and correction beyond threshold. You can click and drag in the plot area to zoom in and hold the **shift** key to pan. Click the **Reset zoom** button to reset zoom.

By default, the graph is displayed for the threshold value of 500 nanoseconds (ns). You can also display data based on a specific threshold value. In the **Threshold (ns)** field, enter the required value in nanoseconds and click **Apply**. Note that the threshold value is persistent in the Nexus Dashboard Fabric Controller settings, and it is used to generate PTP correction threshold Kafka notifications.

In the **Date** field, you can select the appropriate date to view the data. The PTP data is stored up to the last seven (7) days. The default value for the stored data is 7 days. To change this value, navigate to **Settings > Server Settings > IPFM** and set the updated value for the **IPFM history retention in days** field.

In the **Period** field, you can also select a timeframe over which the data has to be displayed. The values you can choose in the **Period** field are Hour (1 hour), 6 hours, 12 hours, or Day (24 hours).

Note that you can click the legends in the graph to hide or display statistics.

If there are any corrections, you can view them in a tabular format by clicking the **Corrections Beyond Threshold** link.

To perform a refresh, click the **Refresh** icon.

### Clock and Port Status

The **Clock & Port Status** tab displays status for Parent Clock, Grandmaster Clock, and Port Status.

The **Port Status** table displays the status of the ports. Click on the **Filter by attributes** field and choose the required attribute, and enter a criteria to filter the port status and press ENTER.

# Policies

Nexus Dashboard Fabric Controller provides the ability to group a set of switches, and allows you to push a set of underlay configurations to the group.

Choose **LAN > Policies** to display the list of policies.

The following table describes the fields that appear on **Fabric Overview > Policies**.

| Field | Description |
|---|---|
| Policy ID | Specifies the policy ID. |
| Switch | Specifies the switch name. |
| IP Address | Specifies the IP address of the switch. |
| Template | Specifies the name of the template. |
| Description | Specifies the description. |
| Entity Name | Specifies the entity name. |
| Entity Type | Specifies the entity type. |
| Source | Specifies the source. |
| Priority | Specifies the priority. |
| Content Type | Species for the content type. |
| Fabric Name | Specifies the fabric name. |
| Serial Number | Specifies the serial number of the switch. |
| Editable | Specifies a Boolean value to indicate if the policy is editable. |
| Mark Deleted | Specifies a Boolean value to indicate if the policy is marked to be deleted. |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **LAN > Policies**.

| Action Item | Description |
|---|---|
| Add Policy | |

| Action Item | Description |
|---|---|
| Edit Policy | Choose a policy from the table and choose **Edit Policy** to modify the policy.<br>**Note**<br>• The policies in the italics font cannot be edited. The value under the **Editable** and **Mark Deleted** columns for these policies is false.<br>• A warning appears when you edit a policy whose **Mark Deleted** value is set to *true*. The switch freeform child policies of **Mark Deleted** policies appears in the Policies dialog box. You can edit only **Python** switch_freeform policies. You cannot edit **Template_CLI** switch_freeform_config policies. |
| Delete Policy | Choose policies from the table and choose **Delete Policy** to delete the policies.<br>**Note**<br>A warning appears when you delete policies whose **Mark Deleted** values are set to *true*.<br>When you delete the policies whose **Mark Deleted** values are set to *true*, these entries are deleted from the NDFC database only but the configs are not deployed to the switch. |
| Generated Config | Choose policies from the table and choose **Generated Config** to view the delta of configuration changes made by every user. |
| Push Config | Choose policies from the table and choose **Push Config** to push the policy configuration to the device.<br>**Note**<br>• This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.<br>• A warning appears if you push configuration for a Python policy.<br>• A warning appears when you push configurations for policies whose **Mark Deleted** values are set to *true*. |

## Adding a Policy

To add a policy, perform the following steps:

### Procedure

**Step 1** Choose **Actions** > **Add Policy**.

The **Create Policy** window appears.

**Step 2** Click and choose required switch and click **Select**.

You must deploy the switch in a pending state.

**Step 3** Click **Choose Template** and choose appropriate policy template and click **Select**.

From Cisco NDFC Release 12.1.2e, you can enable or disable PTP high-correction notification when the system encounters a high-correction event. Whenever the correction value exceeds the configured value then that correction is called a high-correction. By default, a high-correction notification is disabled. Enable it manually to generate the notification. Perform the following steps to enable the high-correction notification:

a. **Enable PTP Telemetry** – Check this check box to enable telemetry for PTP.

b. **Is Large-Scale Fabric?** – Check this check box to generate the high-correction notification. Are there more than 35 devices in the fabric. If yes, PTP events will be used if the switch version is 9.3(5) or higher, or else PTP correction data will be pushed periodically.

c. **PTP High-Correction Interval** – Specify the wait time between two successive notifications, duration value is in seconds.

d. **PTP Correction Range** – Set correction range threshold value (ns), default is 100000 (100us).

From Cisco NDFC Release 12.1.2e, new templates **ipv4_prefix_list** and **ipv6_prefix_list** are added to the template list.

**Step 4** Enter the required name in the **Prefix List Name** field. Perform the following steps to include the prefix-list entries:

a. On the **Prefix-list Entries** field, click **Actions** > **Add**.

The **Add Item** window appears.

b. The mandatory fields on the **Add Item** window are:

**IPv4 Prefix** – Enter the ipv4 prefix address.

**Sequence Number** – Enter the value in the sequence number.

**Action** – From the drop-down list, choose **permit** or **deny**.

Click **Save**.

**Step 5** Repeat the step (5) to add the required number of prefix-list entries.

**Note**
The value in the **Sequence Number** must be higher than the previous prefix-list entry. If not, an error message is displayed.

**Step 6** (Optional) Select the required prefix-list entry and click **Actions > Edit** to edit the selected prefix-list entry.

**Step 7** (Optional) Select the appropriate prefix-list entry and click **Actions > Insert Above** to insert a new prefix-list entry.

**Note**
The value in the **Sequence Number** must be lower than the below prefix-list entry. If not, an error message is displayed.

**Step 8** 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.

**Step 9** Enter the mandatory parameters in the text field and click **Save**.

# Event Analytics

Event Analytics includes the following topics:

# History

The history tab displays information about the deployment and policy change history. Choose **LAN > Fabrics**. Double-click a fabric name to open the **Fabric Overview** window and then click the **History** tab.

# Resources

Cisco Nexus Dashboard Fabric Controller allows you to manage the resources. The following table describes the fields that appear on this page.

| Field | Description |
|---|---|
| Scope Type | Specifies the scope level at which the resources are managed. The scope types can be **Fabric**, **Device**, **Device Interface**, **Device Pair**, and **Link**. |
| Scope | Specifies the resource usage scope. Valid values are the switch serial numbers or fabric names. Resources with serial numbers are unique and can be used on the serial number of the switch only. |
| Device Name | Specifies the name of the device. |
| Device IP | Specifies the IP address of the device. |
| Allocated Resource | Specifies if the resources are managed with device, device interface, or fabric. Valid values are ID type, subnet, or IP addresses. |
| Allocated To | Specifies the entity name for which the resource is allocated. |
| Resource Type | Specifies the resource type. The valid values are **TOP_DOWN_VRF_LAN**, **TOP_DOWN_NETWORK_VLAN**, **LOOPBACK_ID**, **VPC_ID**, and so on. |

| Field | Description |
|---|---|
| Is Allocated? | Specifies if the resource is allocated or not. The value is set to **True** if the resource is permanently allocated to the given entity. The value is set to **False** if the resource is reserved for an entity and not permanently allocated. |
| Allocated On | Specifies the date and time of the resource allocation. |
| ID | Specifies the ID. |

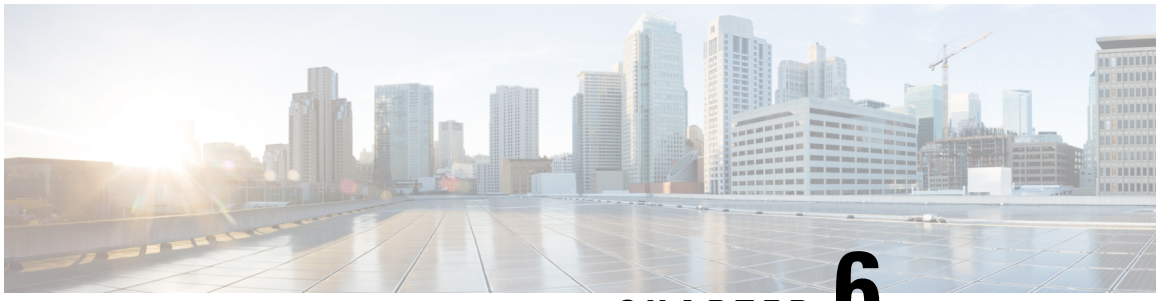
# L4-L7 Services Configuration

Cisco Nexus Dashboard Fabric Controller 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 L4-L7 service devices. You can add a L4-L7 service node, create route peering between the L4-L7 service node and the L4-L7 service leaf switch, and then selectively redirect traffic to these L4-L7 service nodes.



**Note** This is a preview feature in Nexus Dashboard Fabric Controller, Release 12.0.1a. We recommend that you use this feature marked as BETA in your lab setup only. Do not use these features in your production deployment.

From NDFC Release 12.0.2a, you can use this feature in your production environment.



CHAPTER 6

# Policies


- Viewing and Editing Policies, on page 313
- Adding a Policy, on page 315


## Viewing and Editing Policies

Nexus Dashboard Fabric Controller provides the ability to group a set of switches, and allows you to push a set of underlay configurations to the group.

Choose **LAN** > **Policies** to display the list of policies.

The following table describes the fields that appear on **LAN** > **Policies**.

| Field | Description |
|---|---|
| Policy ID | Specifies the policy ID. |
| Switch | Specifies the switch name. |
| IP Address | Specifies the IP address of the switch. |
| Template | Specifies the name of the template. |
| Description | Specifies the description. |
| Entity Name | Specifies the entity name. |
| Entity Type | Specifies the entity type. |
| Source | Specifies the source. |
| Priority | Specifies the priority. |
| Content Type | Species for the content type. |
| Fabric Name | Specifies the fabric name. |
| Serial Number | Specifies the serial number of the switch. |
| Editable | Specifies a Boolean value to indicate if the policy is editable. |

| Field | Description |
|---|---|
| Mark Deleted | Specifies a Boolean value to indicate if the policy is marked to be deleted. |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **LAN > Policies**.

| Action Item | Description |
|---|---|
| Add Policy | To add a policy, see Add Policy |
| Edit Policy | Choose a policy from the table and choose **Edit Policy** to modify the policy.<br>**Note**<br>• The policies in the italics font cannot be edited. The value under the **Editable** and **Mark Deleted** columns for these policies is false.<br>• A warning appears when you edit a policy whose **Mark Deleted** value is set to *true*. The switch freeform child policies of **Mark Deleted** policies appears in the Policies dialog box. You can edit only **Python** switch_freeform policies. You cannot edit **Template_CLI** switch_freeform_config policies. |
| Delete Policy | Choose policies from the table and choose **Delete Policy** to delete the policies.<br>**Note**<br>A warning appears when you delete policies whose **Mark Deleted** values are set to *true*. |
| Generated Config | Choose policies from the table and choose **Generated Config** to view the delta of configuration changes made by every user. |
| Push Config | Choose policies from the table and choose **Push Config** to push the policy configuration to the device.<br>**Note**<br>• This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.<br>• A warning appears if you push configuration for a Python policy.<br>• A warning appears when you push configurations for policies whose **Mark Deleted** values are set to *true*. |

# Adding a Policy

To add a policy, perform the following steps:

**Procedure**

**Step 1** Choose **Actions** > **Add Policy**.

The **Create Policy** window appears.

**Step 2** Click and choose required switch and click **Select**.

You must deploy the switch in a pending state.

**Step 3** Click **Choose Template** and choose appropriate policy template and click **Select**.

From Cisco NDFC Release 12.1.2e, you can enable or disable PTP high-correction notification when the system encounters a high-correction event. Whenever the correction value exceeds the configured value then that correction is called a high-correction. By default, a high-correction notification is disabled. Enable it manually to generate the notification. Perform the following steps to enable the high-correction notification:

a. **Enable PTP Telemetry** – Check this check box to enable telemetry for PTP.

b. **Is Large-Scale Fabric?** – Check this check box to generate the high-correction notification. Are there more than 35 devices in the fabric. If yes, PTP events will be used if the switch version is 9.3(5) or higher, or else PTP correction data will be pushed periodically.

c. **PTP High-Correction Interval** – Specify the wait time between two successive notifications, duration value is in seconds.

d. **PTP Correction Range** – Set correction range threshold value (ns), default is 100000 (100us).

From Cisco NDFC Release 12.1.2e, new templates **ipv4_prefix_list** and **ipv6_prefix_list** are added to the template list.

**Step 4** Enter the required name in the **Prefix List Name** field. Perform the following steps to include the prefix-list entries:

a. On the **Prefix-list Entries** field, click **Actions** > **Add**.

The **Add Item** window appears.

b. The mandatory fields on the **Add Item** window are:

**IPv4 Prefix** – Enter the ipv4 prefix address.

**Sequence Number** – Enter the value in the sequence number.

**Action** – From the drop-down list, choose **permit** or **deny**.

Click **Save**.

**Step 5** Repeat the step (5) to add the required number of prefix-list entries.

**Note**

The value in the **Sequence Number** must be higher than the previous prefix-list entry. If not, an error message is displayed.

**Step 6** (Optional) Select the required prefix-list entry and click **Actions > Edit** to edit the selected prefix-list entry.

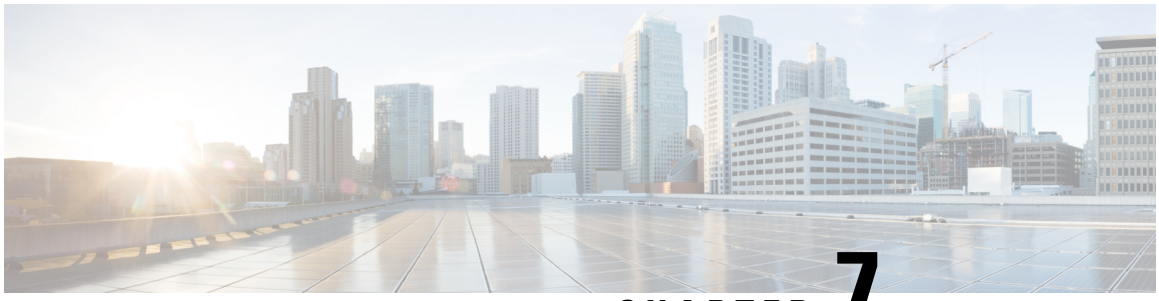
**Step 7** (Optional) Select the appropriate prefix-list entry and click **Actions > Insert Above** to insert a new prefix-list entry.

**Note**
The value in the **Sequence Number** must be lower than the below prefix-list entry. If not, an error message is displayed.

**Step 8** 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.

**Step 9** Enter the mandatory parameters in the text field and click **Save**.



CHAPTER 7

# Interfaces

This section contains the following topics:


- Interfaces, on page 317
- Interface Groups, on page 330



- Interfaces, on page 317
- Interface Groups, on page 330


# 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.

Invalid interface error appears on the following scenarios:

- Interface Mode 'routed' is invalid. Allowed mode is trunk & access.
- Access port which is already allocated to other network.
- Interface which is not available in switch.

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 326.
- 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, int_access_host, 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.

The **Status** column displays the following status of an interface:

- Blue: Pending
- Green: In Sync/Success
- Red: Out-of-Sync/Failed
- Yellow: In Progress
- Grey: Unknown/NA

- If an interface is created out-of-band, you need to perform fabric resync or wait for Config Compliance poling before this interface can be deleted. Otherwise, Config Compliance does not generate the correct diff.

However, you cannot add or edit interfaces for ASR 9000 Series Routers and Arista switches.

You can filter and view information for any of the given fields (such as Device Name). The following table describes the buttons that appear on this page.

**Note**

- Ensure that appropriate configurations are deployed on the Fabric before deploying from the Interfaces option, including proper vPC pair configurations. If you add or edit an interface before configurations are deployed on the Fabric, the configuration may fail on the device.
- Deploy any underlays including vPC Pairing in the fabric before deploying any configurations from the interface manager.

| Field | Description |
|---|---|
| Create Interface | Allows you to add a logical interface such as a port channel, vPC, Straight-through FEX, Active-Active FEX, and loopback.<br>For more information, see Adding Interfaces, on page 323. |
| Create Subinterface | Allows you to add a logical subinterface. |
| Edit interface | Allows you to edit and change policies that are associated with an interface.<br>**Note**<br>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 and IPFM fabrics. |
| Preview interfaces | Allows you to preview the interface configuration. |
| Deploy interfaces | Allows you to deploy or redeploy saved interface configurations. |
| No Shutdown | Allows you to enable an interface (no shutdown or admin up). |
| Shutdown | Allows you to shut down the interface. |
| Add to interface group | Allows you to add an interface to an interface group. |
| Remove from interface group | Allows you to remove an interface from an interface group. |
| Breakout | Allows you to *breakout* an interface. |
| Un-Breakout | Allows you to unbreakout interfaces that are in *breakout* state. |
| Rediscover Interface | Allows you to rediscover or recalculate the compliance status on the selected interfaces. |
| Show commands | Allows you to display the interface show commands. A show command requires show templates in the template library. |
| Deployer History | Allows you to display the interface deployment history details. |

| Field | Description |
|---|---|
| Import | Allows you to import the edited interfaces. The following are the limitations during importing the interfaces:<br>• You are not allowed to import interfaces with the below policy templates:<br>  • All fabric templates with **int_fabric** or **int_ipfm_fabric**.<br>  • **int_vpc_peer** and i**nt_vpc_leaf_tor_assoc**.<br>  • **int_freeform templates**.<br>• You must update the mandatory fields fabric name, serial number, interface name, and policy name.<br>• You are not allowed to import the interfaces with the interface name *nve* and *vlan* except *int_ipfm_vlan* policy. You can import the interface with *int_ipfm_vlan* policy.<br>• The allowed MTU range for integer values is between 576 and 9216.<br>• The allowed MTU string values is either *default* or *jumbo*.<br>• The fabric name, serial number, and interface name must be unique.<br>• Single policy import per CSV is only allowed and can not combine with multiple policy types.<br>**Note**<br>There is a server property to set the maximum number of rows that can be imported. By default, the property is 200 for import. |
| Export | Allows you to export the selected interfaces with multiple types of policies.<br>**Note**<br>There is no limit for number of interfaces to export. There is a server property to set the maximum number of rows that can be exported. By default, the property is 200 for export. |
| Delete Interface | 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. |

The following table describes the new user role access-admin operations support in the host facing port of **Interfaces** window from Cisco Nexus Dashboard Fabric Controller Release 11.5(1).

| Operations | User Roles |
|---|---|
| | access-admin |
| Create new interface | Save, Preview, Deploy |
| Breakout | Blocked |
| Un-Breakout | Blocked |
| Edit interface | Save, Deploy |
| Delete Interface | Save, Deploy |
| Shutdown | Save, Deploy |
| No Shutdown | Save, Deploy |
| Show commands | Clear Output, Execute |
| Rediscover interface | Supported |
| Deploy Interfaces | Cancel, Deploy Config |
| Import | Supported |
| Export | Supported |

You can disable deployments, or freeze, a fabric in Nexus Dashboard Fabric Controller 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 | Nexus Dashboard Fabric Controller 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 |

| Operations | Nexus Dashboard Fabric Controller Mode | |
|---|---|---|
| | Freeze Mode | Monitor Mode |
| Rediscover | Supported | Supported |
| Deploy | Blocked | Blocked |
| Import | Supported | Supported |
| Export | Supported | Supported |

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**.

The following table describes the fields that appear on **LAN > Interfaces > Interfaces**.

| Field | Description |
|---|---|
| Fabric Name | Specifies the fabric name. |
| Device Name | Specifies the device name. |
| Interface | Specifies the interface name. |
| Admin Status | Specifies the administrative status of the interface. The status can be either Up or Down. |

| Field | Description |
|---|---|
| Oper-Status | Specifies the operational status of the interface. The status can be either Up or Down. |
| Reason | Specifies the reason. |
| Policies | Specifies the policy name. |
| Overlay Network | Specifies the overlay network. |
| Sync Status | Specifies the sync status. Specifies if the interface status is In-Sync or Out-Of-Sync. |
| Interface Group | Specifies the interface group to which the interface belongs to. |
| Port Channel ID | Specified the port channel ID. |
| vPC ID | Specifies the vPC ID. |
| Speed | Specifies the interface speed. |
| MTU | Specifies the MTU size. |
| Mode | Specifies the interface mode. |
| VLANs | Specifies the VLANs. |
| IP/Prefix | Specifies the interface IP/Prefix. |
| VRF | Specifies virtual routing and forwarding instances (VRFs). |
| Neighbour | Specifies the interface neighbour. |
| Description | Specifies the interface description.<br>**Note**<br>If the interface description is more than 64 characters, you must configure the switch using **snmp ifmib ifalias long** command. |

# Adding Interfaces

To add the interfaces from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN > Interfaces > Interfaces**.

**Step 2** Click **Actions > Create new interface** to add a logical interface.

The **Create new interface** window appears.

**Step 3** From 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 Nexus Dashboard Fabric Controller, 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 and deploy vPC and peer-link configurations using the **Save Deploy** option. After 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** 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. .

**Step 4** In the **Select a device** field, choose a device.

Devices are listed based on the fabric and interface type.. 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**, **Tunnel ID**, **Interface name**, **VLAN 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** Under the **Policy** field, select a 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 **Policy Options**.

The fields vary according to the interface type you choose.

**Note**
From Cisco Nexus Dashboard Fabric Controller Release 11.5(1) you can mirror the configurations of Peer-1 on Peer-2 while creating a vPC. When you check the **Enable Config Mirroring** check box, the Peer-2 fields will be grayed out. The configurations that you enter in the Peer-1 fields will be copied to Peer-2 fields.

From Cisco NDFC Release , you can set Native Vlan for the interface which has **int_trunk_host** or **int_port_channel_trunk_host**, or **int_vpc_trunk_host** policy template.

A trunk port can carry nontagged packets simultaneously with tagged packets. When you assign a default port VLAN ID to the trunk port, all untagged traffic travels on the default port VLAN ID for the trunk port, and all untagged traffic is assumed to belong to this VLAN. This VLAN is referred to as the native VLAN ID for a trunk port. That is, the native VLAN ID is the VLAN that carries untagged traffic on trunk ports.

**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** and **Un-Breakout**: You can break out and unbreakout an interface by using the **Breakout** and **Un-Breakout** options.

# Breakout

To breakout an interface, from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. On Interface window, click **Actions** > **Breakout**.

   The **Breakout Interfaces** window appears.

2. Choose the required option in the window and click **Breakout**.

   The available options are 10g-4x, 25g-4x, 50g-2x, 50g-4x, 100g-2x, 100g-4x, 200g-2x, and Unbreakout.

# UnBreakout

You can unbreakout interface that are in breakout state.

On Interface window, click **Actions** > **UnBreakout**.

**Note** The Interface which are not in breakout state, the unbreakout option is grayed out.

# Editing Interfaces

To edit the interfaces from the Cisco Nexus Dashboard Fabric Controller 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 **LAN > Interfaces > Interfaces**.

You can break out and unbreak out an interface by using the breakout option in the **Actions** menu.

**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 **Actions > Edit interface** to edit an interface.

The variables that are shown in the **Edit interface** window are based on the template and its policy. Select the appropriate policy. Save the policy 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** tab. You can update such interfaces.

For interface policies that are not created from the **LAN > Interfaces > 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 policy is used to create a loopback interface. You can edit the loopback IP address and description but not the int_fabric_loopback policy instance.
- Fabric underlay network interface policies (int_fabric_num, 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 **LAN > Interfaces > Interfaces**.

**Step 2** Select a link and click **Actions > More > Rediscover Interface**.

# Deleting Interfaces

To delete the interfaces from the Cisco Nexus Dashboard Fabric Controller 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 **LAN > Interfaces > Interfaces**.

**Step 2** Select the interfaces.

**Step 3** Click **Actions > More > Delete Interface**.

You cannot delete logical interfaces created in the fabric underlay.

**Step 4** Click **Save**.

**Step 5** Click **Deploy** to delete the interface.

# Shutting Down and Bringing Up Interfaces

To shut down and bring up the interfaces from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN > Interfaces > Interfaces**.

**Step 2** Select the interfaces that you want to shut down or bring up.

**Step 3** Click **Shutdown** to disable the selected interfaces. For example, you may want to isolate a host from the network or a host that is not active in the network.

A confirmation window appears where you can save, preview, and deploy the changes. Click **Save** to preview of deploy the changes.

**Step 4** Click **No Shutdown** to bring up the selected interfaces.

A confirmation window appears where you can save, preview, and deploy the changes. Click **Save** to preview or deploy the changes.

# Viewing Interface Configuration

To view the interface configuration commands and execute them from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN > Interfaces > Interfaces**.

Select the interface whose configurations you want to view and click **Actions > More > Show commands**.

**Step 2** In the **Interface show commands** window, select the action from the **Commands** drop-down box and click **Execute**. The interface configurations are displayed on 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 **Templates**.

# Rediscovering Interfaces

To rediscover the interfaces from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN > Interfaces > Interfaces**.

**Step 2** Select the interfaces that you want to rediscover and click **Actions > More > Rediscover Interface** 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 Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN > Interfaces > Interfaces**.

**Step 2** Select the interface and click **Actions > More > Deployer History** to view the configuration history on the interface.

**Step 3** Click **Status** to view each command that is configured for that configuration instance.

# Deploying Interface Configurations

To deploy the interface configuration from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN > Interfaces > Interfaces**.

**Step 2** Select an interface that you want to deploy and click **Actions > Deploy Interfaces** to deploy or redeploy configurations that are saved for the interface.

**Note**
You can select multiple interfaces and deploy pending configurations.

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 deploy for the switch
- During a deploy
- During hourly sync

# Creating External Fabric Interfaces

You can add and edit port channel, vPC, subinterface, and loopback interfaces for external fabric devices. You cannot add Straight-through FEX and Active-Active FEX functions.

The Breakout port function is only supported for the Cisco Nexus 9000, 3000, and 7000 Series Switches in the external fabric.

When you add an interface to an external fabric device, the Resource Manager is not in sync with the device. So, ensure that the value populated in the ID field (Port-channel ID, vPC ID, Loopback ID, etc) is not previously configured on the switch.

If you want to configure a portchannel in the external fabric, you should add and deploy the **feature_lacp** policy on the switches where the portchannel will be configured.

When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches. If you click **Save & Deploy** in the fabric topology screen, it displays an error message. However, the following settings (available when you right-click the switch icon) are allowed:

vPC pairing - You can designate a vPC switch pair, but it is only for reference.

View/edit policy - You can add a policy but you cannot deploy it on the switch.

Manage interfaces – You can only create intent for adding interfaces. If you try to deploy, edit, or delete interfaces, it results in an error message.

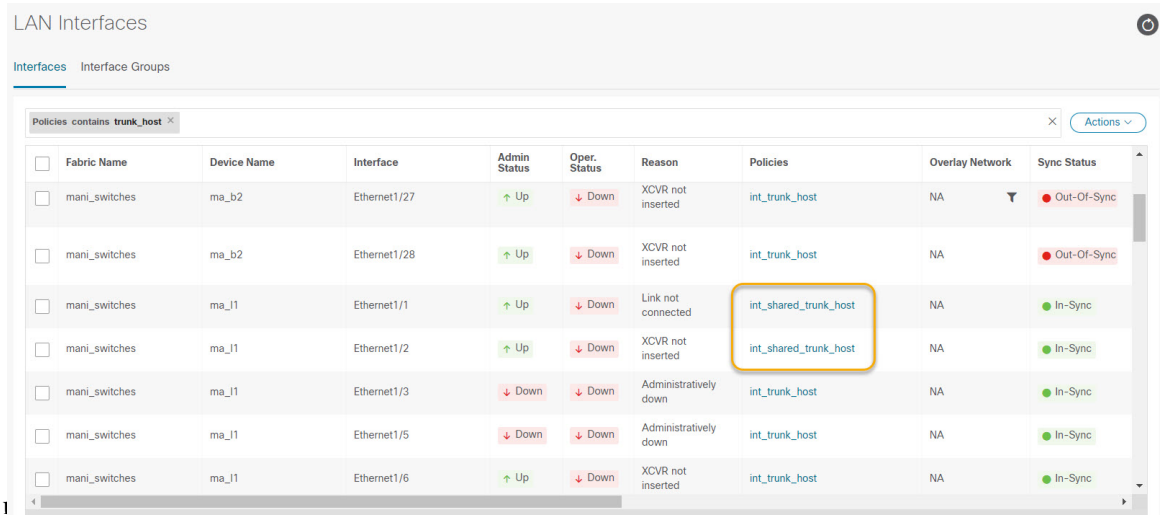
# Interface Groups

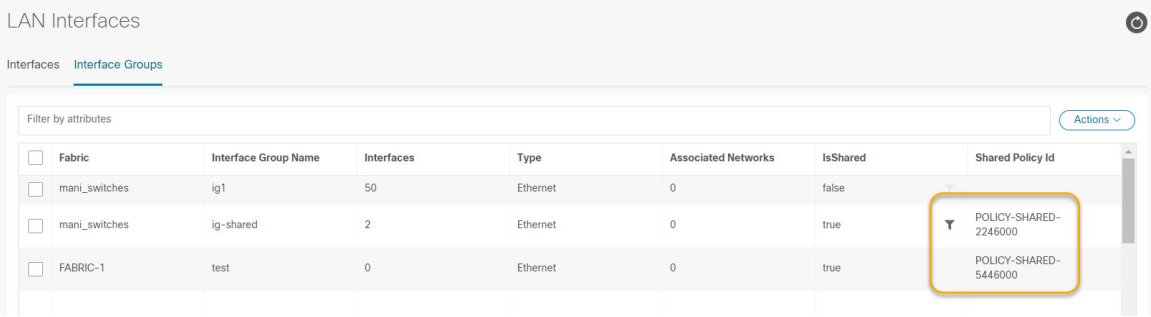
An interface group consists of multiple interfaces with same attributes. You can create an interface group that allows grouping of host-facing interfaces at a fabric level. Specifically, you can create an interface group for physical Ethernet interfaces, Layer 2 port-channels, and vPCs. You can attach or detach multiple overlay networks to the interfaces in an interface group.

### Shared Policy

From Cisco NDFC Release 12.1.2e, you can create and add a shared policy for an interface group. This policy ensure update of appropriate configurations for all the interfaces in the Interface group. In the shared policy, all the interfaces will have the same underlay and overlay attributes. When you change the configuration in the shared policy, then that configuration is applied to all the interfaces.

You can see the details of shared policy under the **Policies**






Custom policy can also be created by selecting the policy from the template list and **Duplicate Template** to add the additional information. The shared policy must contain the tags **interface_edit_policy**, **interface_edit_shared_policy**, and **int_trunk**.

### Guidelines

- Interface groups are only supported for the fabrics with the **Data Center VXLAN EVPN** 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 Layer 2 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.
- From Cisco NDFC Release 12.1.2e, you can include Layer 2 ToR interfaces in the interface groups.
- Interfaces added to an interface group with a shared policy replaces the individual policy and get the shared policy.
- You can change description and status of each interface in interface group.
- Interface removed from an interface group with a shared policy will set to a default policy.
- VMs should have the same configuration for all the interfaces under the shared policy.
- Shared policy is supported only for **Ethernet** interface type in interface group.
- Ethernet interface groups now have a common policy.
- **Port-Channel** and **vPC** interface types are not supported for adding shared policy in interface group.
- When the MTU value in the shared policy has to be changed, make sure to update the fabric settings with the same value across all switches of that fabric.
- When the **ptp**, **ttag**, and **ttag-strip** option from the shared policy has to be used, make sure to enable PTP globally in fabric settings.
- When the netflow option has to be used, make sure all the interfaces of interface groups are capable of netflow configuration and that it is enabled globally in fabric settings.

- For Layer 2 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 **Add to Interface Group** and **Remove from 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 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

To create an interface group from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN > Interfaces > Interface Groups**.

**Step 2** Click **Actions > Create new interface group**.

**Step 3** From the **Select Fabric** window, select a fabric and click **Select**.

**Step 4** In the **Create new interface group** window, provide an interface group name in the **Interface Group Name** field, select an Interface Type, and click **Save**.

An interface group name can have a maximum length of 64 characters

**Note**
An interface can be added to single interface group only.

**Step 5** Click on the **Policy** field, the **Select Policy** window appears. Choose **int_shared_trunk_host** policy and then click **Select**.

From Cisco NDFC Release 12.1.2e, a new **Policy** field is introduced in the **Create new interface group** window. You can add a shared policy to the interface group which can be shared by the interfaces existing in that group. Shared Policy is optional, for upgrades, all the existing interface group will not have a policy.

.

**Note**
The policy field supports the **Ethernet** interface type only.

**Step 6** Enter the mandatory parameters in the text field and click **Save**.

**Step 7** Click the **Interfaces** tab.

**Step 8** Select the interfaces that have to be grouped and click **Actions > Add to interface Group**.

**Step 9** In the **Add new interface Group** window, create a custom interface group by entering an interface group name in the **Select Interface Group** field and click **Create custom** .

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.

You can create interface groups from either the **Interfaces Groups** window or the **Interfaces** window under Fabric Overview.

**Step 10** Click **Save**.

In the **Interfaces** window, you can see the interface group name under the **Interface Group** column.

**Step 11** To edit an interface group, click **Actions** > **Edit Interface Group**. You can update the policy options after you assigned the shared policy.

**Note**
You cannot edit or delete the shared policy template.

# Removing Interfaces from an Interface Group

To remove interfaces from an interface group from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN > Interfaces**.

**Step 2** Select the interfaces to disassociate from an interface group and click **Actions > Remove from interface Group**.

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

To attach networks to an interface group from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Double click on the fabric to launch **Fabric Overview**.

**Step 2** On the **Networks** tab, 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 3** In the **Interface Groups** window, you can perform the following:

- Select an existing interface group from the **Select Interface Group** drop-down list and click **Save**.

  For example, you select three networks and the interface group **test**, and click the **Save** button, the following operations are performed in the background:

  **a.** Nexus Dashboard Fabric Controller retrieves interfaces that are part of the interface group **test**.

  **b.** Nexus Dashboard Fabric Controller 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, Nexus Dashboard Fabric Controller pushes the "**switchport trunk allowed vlan add xxxx**" command three times for each selected network.

  **Note**
  Nexus Dashboard Fabric Controller ensures that there's no duplicate configuration intent.

  If you click the **Clear** button, Nexus Dashboard Fabric Controller 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 new interface group**. Click **Save**.

  If you choose this option, make sure to add interfaces to this Interface Group in the **Interfaces** window. As a result, Nexus Dashboard Fabric Controller 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 332.

**Step 4** Click **Actions > Recalculate & Deploy** to deploy the selected networks on the switches.

# Detaching a Network from an Interface Group

This procedure shows how to detach a network from an interface group in the Networks window. Also, you can detach 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** Double click on the fabric to launch **Fabric Overview**.

**Step 2** On the **Networks** tab, select the networks that you need to detach to an interface group and click **Add to Interface Group**.

**Step 3** In the **Add to Interface Groups** window, select the interface group from the **Select Interface Group** drop-down list and click **Clear** to detach a network.

**Step 4** (Optional) Navigate to **LAN > Interfaces**.

Under the **Overlay Network** column, you can see the detached network in the red color for the corresponding interface. Click the network to view the expected config that is struck through.

**Step 5** Navigate to the **Networks** tab and choose **Actions > Recalculate & Deploy**.

# Deleting an Interface Group

An interface group is automatically deleted when it's not in use. Nexus Dashboard Fabric Controller performs an implicit delete of an interface group if there are no interfaces and no networks mapped to the interface group. You can perform an explicit delete by clicking on **Interface Group > Actions > Delete 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.

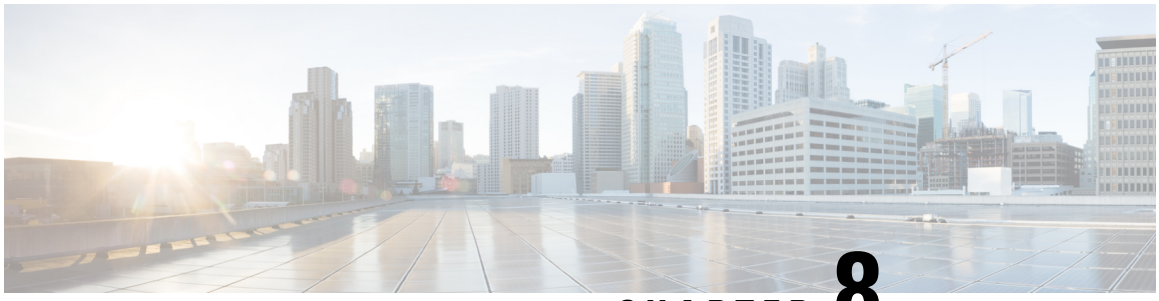
**Step 4** Click **Clear**.



PART II

# Virtual Management


- Virtual Infrastructure Manager, on page 339



CHAPTER 8

# Virtual Infrastructure Manager


- Virtual Infrastructure Manager, on page 339
- **Adding vCenter Visualization**, on page 342
- Kubernetes Cluster, on page 344
- OpenStack Cluster, on page 347
- Annexure, on page 349


## Virtual Infrastructure Manager

UI Path: **Virtual Management** > **Virtual Infrastructure Manager**

**Note** Ensure that you have enabled Network visualization of Virtual Machines feature for Cisco Nexus Dashboard Fabric Controller.

1. Choose **Settings** > **Feature Management**, choose the following check boxes:
   - Kubernetes Visualizer
   - VMM Visualizer
   - Openstack Visualizer
2. Click **Apply**.

The following table describes the fields that appear on Virtual Infrastructure Manager window:

| Field | Description |
|---|---|
| Server | Specifies the Server IP Address. |
| Type | Specifies the type of instance that can be one of the following:<br>• vCenter<br>• Kubernetes Cluster<br>• OpenStack Cluster |

| Field | Description |
|---|---|
| Managed | Specifies the status of the cluster either Managed or Unmanaged. |
| Status | Specifies the status of the added cluster. |
| User | Specifies the user created the cluster. |
| LastUpdated Time | Specifies the last updated time for the cluster. |

**Note** Click **Refresh** icon to refresh the Virtual Infrastructure Manager table.

The following table describes the action items, in the Actions menu drop-down list, that appear on Virtual Infrastructure Manager window:

| Action Item | Description |
|---|---|
| Add Instance | From the **Actions** drop-down list, choose **Add Instance**. For more instructions, see Adding an Instance.<br>**Note**<br>Ensure that you have configured same IP address on Routes. Refer to Configuring Routes IP Address. |
| Edit Instance | Choose an instance to edit. From the **Actions** drop-down list, choose **Edit Instance**. Make the necessary changes and click **Save**. Click **Cancel** to discard the changes. |
| Delete Instance(s) | Choose one or more required instance to delete. From the **Actions** drop-down list, choose **Delete Instance(s)**. Click **Confirm** to delete the instance. Click **Cancel** to discard the delete. |
| Rediscover Instance(s) | Choose one or more required instance to rediscover. From the **Actions** drop-down list, choose **Rediscover Instance(s)**. A confirmation message appears. |

For more information:

# Support for Cisco UCS B-Series Blade Servers

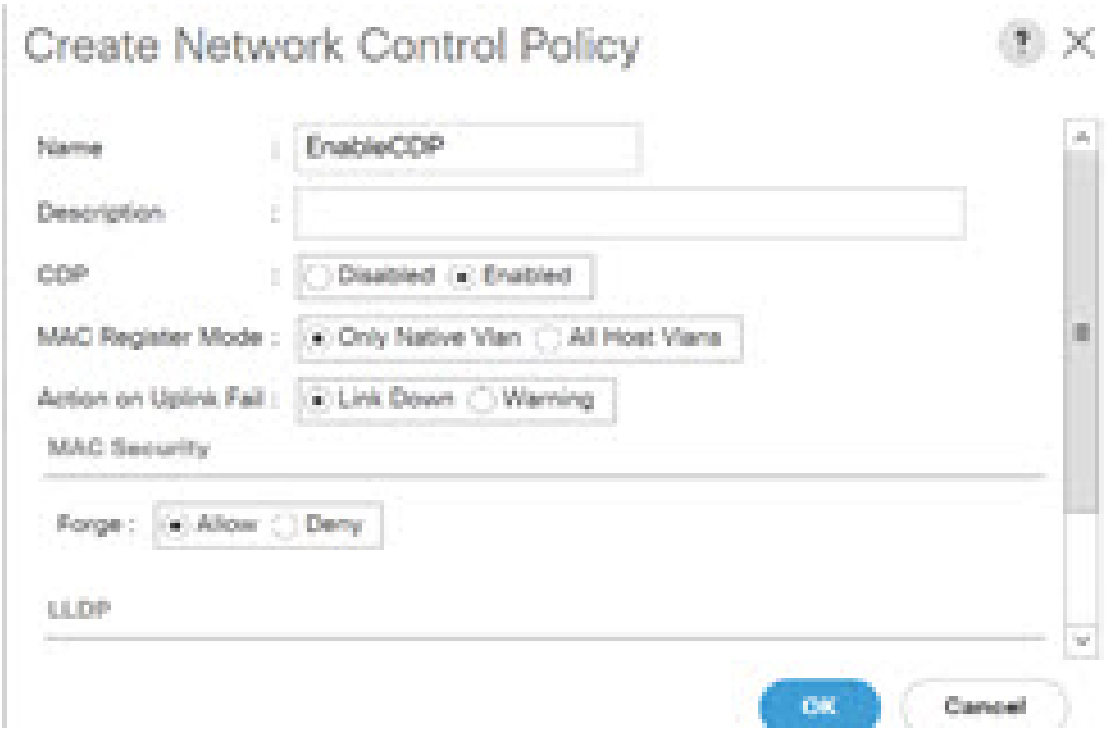
NDFC supports hosts running on UCS type B (chassis UCS) that are behind the Fabric interconnect. You must enable CDP of the vNIC on Cisco UCSM to use this feature.

**Note** By default, CDP is disabled on Cisco UCSM.

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

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

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




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

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

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

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

# Configuring Routes IP Address

Before you add IP address to vCenter, you must configure same IP address on Cisco Nexus Dashboard.

To configure Routes on Cisco Nexus Dashboard, perform the following steps:

**Procedure**

**Step 1** Choose **Infrastructure** > **Cluster Configuration**.

**Step 2** On **General** tab, in **Routes** card, click **Edit** icon.

The **Routes** window appears.

**Step 3** To configure IP addresses, click **Add Management Network Routes**, enter required IP addresses, and click **check** icon.

**Step 4** Click **Save**.

The route configuration is governed by following two scenarios:

**a.** For vCenter, which is an application server is typically reachable over mgmt network.

**b.** The ESXi servers that are managed by vCenters and the baremetal servers hosting the K8s instances and/or OpenStack instances would be connected to the fabric network directly. Hence, they will be reachable over data networks.

# Adding vCenter Visualization

You can perform various actions in the **Actions** menu drop-down list, that appear on **Virtual Management > Virtual Infrastructure Manager.**

**Procedure**

**Step 1** Choose **Actions** > **Add Instance**.

The **Add Instance** window appears.




**Step 2** Choose **vCenter** from Select Type drop-down list.

Enter required IP address or Domain name and password in the respective fields.

**Step 3** Click **Add**.

You can view added vCenter cluster in the **Virtual Infrastructure Manager** window.

**Step 4** To edit an instance**,** choose required vCenter, choose **Actions > Edit Instance** and click **Save** changes.

You can update password for the selected vCenter cluster and change the admin status to Managed or Unmanaged and vice-versa.

**Note**
For the vCenter cluster in Unmanaged status, you cannot view the topology and vCenter cluster details on dashboard.

**Step 5** To delete one or more vCenter cluster, choose the required vCenter, choose **Actions** > **Delete Instance(s)** and click **Confirm** changes.

**Note**
All the data will be deleted if you delete the Cluster. The Cluster will be removed from the Topology view also.

**Step 6** To rediscover one or more vCenter cluster, choose the required vCenter, choose **Actions > Rediscover Instance(s)**.

A confirmation message appears.

# Kubernetes Cluster

**Note** Ensure that you have enabled Network Visualization of K8s clusters feature for Cisco Nexus Dashboard Fabric Controller .

Choose **Settings** > **Feature Management** choose check box **Kubernetes Visualizer** and click **Apply**.

You can view the added Kubernetes Visualizer details on dashboard. Navigate **Dashboard** > **Kubernetes Pods**

To enable LLDP on NDFC, choose **Settings** > **Server** > **Settings** > **Discovery**. Choose check box **enable / disable neighbor link discovery using LLDP**.

**Note** LLDP is applicable for Bare-metal Kubernetes clusters only.

- Ensure that the LLDP feature is enabled on all fabric switches for which the cluster node is connected. (Switches may be spine or leaf switches).
- 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.

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.

**Note** When discovering or visualizing VM-based Kubernetes cluster, it must first onboard the vCenter cluster which is managing the VMs hosting the Kubernetes cluster being discovered. Without this, Kubernetes cluster discovery would result in failure.

# Configuring Routes IP Address

Before you add IP address to Kubernetes cluster, you must configure same IP address on Cisco Nexus Dashboard.

To configure Routes on Cisco Nexus Dashboard, perform the following steps:

**Procedure**

**Step 1** Choose **Infrastructure** > **Cluster Configuration**.

**Step 2** On **General** tab, in **Routes** card, click **Edit** icon.

The **Routes** window appears.

**Step 3** To configure IP addresses, click **Add Management Network Routes**, enter required IP addresses, and click **check** icon.

**Step 4** Click **Save**.

# Adding Kubernetes Cluster

You can perform various actions in the **Actions** menu drop-down list, that appear on **Virtual Management** > **Virtual Infrastructure Manager**.

**Note** Ensure that you have configured same IP address on Routes. Refer to Configuring Routes IP Address.

**Procedure**

**Step 1** Choose **Actions** > **Add Instance**

The **Add Instance** window appears.

**Step 2** Choose **Kubernetes Cluster** from Select Type drop-down list.

**Step 3** Enter **Cluster IP address**, **Username** in appropriate fields.

**Step 4** 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.

Refer to Annexure, on page 349 to generate the certificate **genk8clientcert.sh**.

**Step 5** Login to the Kubernetes cluster controller node.

You need admin privileges to generate the certificates.

**Step 6** Copy the genk8clientcert.sh and kubereader.csr from the NDFC server location to the Kubernetes Cluster controller node.

Perform a "vnc cut and paste" operation to ensure that all the characters are copied correctly.

**Step 7** 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 3).
- 10.x.x.x is the IP address of the NDFC server.

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 NDFC IP address)

```
dcnm(root)# cat k8s_cluster_ca.crt
```

**Step 8** 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 NDFC.

**Step 9** 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 10** Copy the content in the k8s_cluster_ca.crt to **Cluster Certificate** field.

**Note**
Perform a "vnc cut and paste" operation to ensure that all the characters are copied correctly.

**Step 11** Click **Add**.

You can view added Kubernetes cluster in the **Virtual Infrastructure Manager** window.

**Note**
You can view details of the added Kubernetes cluster details on the dashboard and topology window. Navigate **Dashboard** > **Kubernetes Pods** and topology window.

**Step 12** To edit Kubernetes cluster, choose required cluster, choose **Actions** > **Edit Instance**, click Edit to 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.

**Note**
For the kubernetes cluster in Unmanaged status, you cannot view the topology and Kubernetes cluster details on dashboard.

**Step 13** Click **Save** to save the changes or click **Cancel** to discard changes.

**Step 14** To delete one or more Kubernetes Cluster, choose the required cluster, choose **Actions** > **Delete Instance(s)** to delete the cluster.

**Note**
All the data will be deleted if you delete the Cluster. The Cluster will be removed from the Topology view also.

**Step 15** Click **Confirm** to delete the cluster.

**Step 16** To rediscover one or more Kubernetes cluster, choose required Kubernetes cluster, choose **Actions** > **Rediscover Instance(s)**.

A confirmation message appears.

# OpenStack Cluster



**Note** This is a preview feature in Nexus Dashboard Fabric Controller, Release 12.0.2. We recommend that you use this feature marked as BETA in your lab setup only. Do not use these features in your production deployment.



**Note**

- Ensure that you have enabled Network Visualization of Openstack Clusters feature for Cisco Nexus Dashboard Fabric Controller. Choose **Settings > Feature Management** choose check box **Openstack Visualizer** and click **Apply**.
- Ensure that the vCenter cluster or Kubernetes cluster feature must be enabled to add an openstack cluster.

- To enable LLDP on NDFC, choose Web UI, choose **Settings > Server Settings > Discovery.** Choose check box **enable / disable neighbor link discovery using LLDP**.

- 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.
- 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
```

# Configuring Routes IP Address

Before you add IP address to Openstack Visualizer, you must configure same IP address on Cisco Nexus Dashboard.

To configure Routes on Cisco Nexus Dashboard, perform the following steps:

**Procedure**

**Step 1** Choose **Infrastructure** > **Cluster Configuration**.

**Step 2** On **General** tab, in **Routes** card, click **Edit** icon.

The **Routes** window appears.

**Step 3** To configure IP addresses, click **Add Management Network Routes**, enter required IP addresses, and click **check** icon.

**Step 4** Click **Save**.

# Configuring AMQP Endpoints on OpenStack Cluster

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

  Make the following configuration changes in the OpenStack Nova service. Replace the parameter values as shown. The Nova configuration file is located at the path:

  ```
  /etc/nova/nova.conf
  ```

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

  [oslo_messaging_notifications]
  ```

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



**Note**

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

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

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

# Annexure

The following message is displayed, after the certificates are generated successfully:

```
#!/usr/bin/bash
############################################################
# Title: Script to provision the client CSR and generat the #
#        the client SSL certificate.                       #
############################################################

# Create CSR resource template.
function create_csr_resource() {
    K8SUSER=$1
    DCNM=$2
    FILE=${K8SUSER}_${DCNM}_csr_res.yaml
    echo "
apiVersion: certificates.k8s.io/v1
kind: CertificateSigningRequest
metadata:
  name: ${K8SUSER}_${DCNM}csr
spec:
  groups:
  - system:authenticated
  request: ${BASE64_CSR}
  signerName: kubernetes.io/kube-apiserver-client
  usages:
  - digital signature
  - key encipherment
```

```
  - client auth" > $FILE
}

# Create CLUSTER ROLE resource template
function create_cluster_role() {
    K8SUSER=$1
    DCNM=$2
    FILE=${K8SUSER}_${DCNM}_cluster_role_res.yaml
    echo "
kind: ClusterRole
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: clustrole_${K8SUSER}_${DCNM}
rules:
- apiGroups: [\"\"]
  resources: ["nodes", "namespaces", "pods", "services"]
  verbs: ["get", "list", "watch"]" > $FILE
}

# Create CLUSTER ROLE BINDING template
function create_cluster_role_binding() {
    K8SUSER=$1
    DCNM=$2
    FILE=${K8SUSER}_${DCNM}_cluster_rolebinding_res.yaml
    echo "
kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: clustrolebind_${K8SUSER}_${DCNM}
roleRef:
  kind: ClusterRole
  name: clustrole_${K8SUSER}_${DCNM}
  apiGroup: rbac.authorization.k8s.io
subjects:
- kind: User
  name: ${K8SUSER}
  apiGroup: rbac.authorization.k8s.io" > $FILE
}

function valid_ip() {
    local  ip=$1
    local  stat=1

    if [[ $ip =~ ^[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}$ ]]; then
        OIFS=$IFS
        IFS='.'
        ip=($ip)
        IFS=$OIFS
        [[ ${ip[0]} -le 255 && ${ip[1]} -le 255 \
            && ${ip[2]} -le 255 && ${ip[3]} -le 255 ]]
        stat=$?
    fi
    return $stat
}

# Start of the script
if [ "$#" -ne 2 ]; then
    echo "Please provide the username and IP of the DCNM"
    echo
    exit 1
else

    # Check if user have required K8s privileges
    LINUX_USER=$(whoami)
```

```
K8S_CONF_PATH=""
echo
echo "Hello ${LINUX_USER}! I am going to help you generate K8s cluster CA and K8s client certificate."

if [ ${LINUX_USER} == "root" ] ; then
    # You are root
    if [ ! -d "/root/.kube" ] ; then
        echo
        echo "Directory /root/.kube does not exists."
        echo "User ${LINUX_USER} does not have required K8s privileges"
        echo "Please make sure you are logged into K8s cluster's master node"
        echo
        exit 1
    else
        K8S_CONF_PATH=/${LINUX_USER}/.kube/config
    fi
else
    # You are not root
    if [ ! -d "/home/${LINUX_USER}/.kube" ] ; then
        echo
        echo "Directory /home/${LINUX_USER}/.kube does not exists."
        echo "User ${LINUX_USER} does not have required K8s privileges"
        echo "Please make sure you are logged into K8s cluster's master node"
        echo
        exit 1
    else
        K8S_CONF_PATH=/home/${LINUX_USER}/.kube/config
    fi
fi

# Check if K8s config file exist
if [ ! -f ${K8S_CONF_PATH} ]; then
    echo
    echo "${K8S_CONF_PATH} file does not exist"
    echo "K8s CA certificate can not be exported"
    echo "Please make sure you are logged into K8s cluster's master node"
    echo
    exit 1
fi

K8SUSER=$1
DCNM=$2
K8S_CA_CRT="k8s_cluster_ca.crt"

# Validate the IP address
if valid_ip $DCNM; then
    echo -e
else
    echo "${2} is not a valid IP address"
    echo
    exit 1
fi

# Validate the CSR file format
if [ ${K8SUSER: -4} == ".csr" ]; then
    K8SUSER=${K8SUSER%.csr}
fi

if [ ! -f "./${K8SUSER}.csr" ]; then
    echo
    echo "./${K8SUSER}.csr does not exist"
    echo "CSR file is required for creation of client certificate"
    echo
```

```
        exit 1
    fi

    echo "Generating certificate for ${K8SUSER} for DCNM ${DCNM}"
    echo

    # Encoding the .csr file in base64
    export BASE64_CSR=$(cat ./${K8SUSER}.csr | tr -d '\n')

    # Create the CSR resource in K8s cluster
    create_csr_resource $K8SUSER $DCNM

    # Delete if the CSR resource already exist. We need a fresh one.
    kubectl delete csr ${K8SUSER}_${DCNM}csr &> /dev/null
    status=$?
    if test $status -eq 0
    then
        echo "./${K8SUSER}_${DCNM}csr CSR resource already exist, removing it"
    else
        echo "./${K8SUSER}_${DCNM}csr CSR resource does not exist, creating it"
    fi

    # Create the CertificateSigninRequest resource
    kubectl apply -f ${K8SUSER}_${DCNM}_csr_res.yaml

    # Check the status of the newly created CSR
    kubectl get csr

    # Approve this CSR
    echo "Approving the CSR"
    kubectl certificate approve ${K8SUSER}_${DCNM}csr

    # Check the status of the newly created CSR
    kubectl get csr

    # Create role resource definition
    kubectl delete clusterrole clustrole_${K8SUSER}_${DCNM} &> /dev/null
    create_cluster_role $K8SUSER $DCNM
    kubectl apply -f ${K8SUSER}_${DCNM}_cluster_role_res.yaml

    # Create role binding definition
    kubectl delete clusterrolebinding clustrolebind_${K8SUSER}_${DCNM} &> /dev/null
    create_cluster_role_binding $K8SUSER $DCNM
    kubectl apply -f ${K8SUSER}_${DCNM}_cluster_rolebinding_res.yaml

    # Extract the client certificate
    echo "Extracting the user SSL certificate"
    kubectl get csr ${K8SUSER}_${DCNM}csr -o jsonpath='{.status.certificate}' >
${K8SUSER}_${DCNM}.crt
    echo "" >> ${K8SUSER}_${DCNM}.crt

    # Export the K8s cluster CA cert
    if [ -f ${K8S_CONF_PATH} ]; then
        echo "Exporting K8s CA certificate"
        cat ${K8S_CONF_PATH} | grep certificate-authority-data | awk -F ' ' '{print $2}' >
 ${K8S_CA_CRT}
    fi
    echo
    echo "----------------------------------------------------------------"
    echo "Notes: "
    echo "1. The K8s CA certificate is copied into ${K8S_CA_CRT} file."
    echo "      This to be copied into \"Cluster CA\" field."
    echo "2. The client certificate is copied into ${K8SUSER}_${DCNM}.crt file."
    echo "      This to be copied into \"Client Certificate\" field."
```

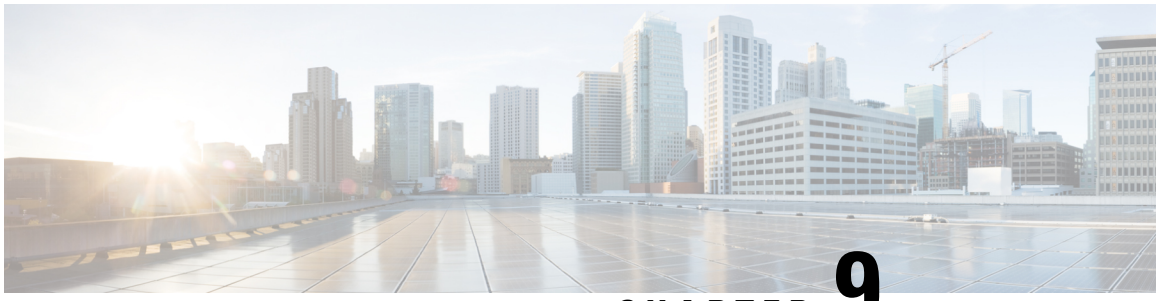
```
    echo "------------------------------------------------------------"
    echo
fi
```



PART III

# Settings


- Server Settings, on page 357
- Feature Management, on page 359
- Credentials Management, on page 363



CHAPTER 9

# Server Settings


• Server Settings, on page 357


# Server Settings

You can set the parameters that are populated as default values.

To set the parameters of the Nexus Dashboard Fabric Controller server from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Choose **Settings** > **Server Settings**.

   Server settings are classified under different tabs,

2. Modify the settings based on the requirement.
3. Click **Save** to apply the new modified settings.

Each microservice of enabled features has other tabs and properties other than listed below. Each field has short description. If there is error during configuring any features, corresponding tab is highlighted in red, and **Save** button is disabled till the errors are resolved. Comprehensive checks are performed in NDFC server by the microservices, if there are any errors is displayed on NDFC UI. Server settings supported for 'all-or-none' to save properties and it doesn't support partial updates.

**Note** You can modify required properties in server settings without support of Cisco TAC.

### LAN Device Management Connectivity under Admin

From Cisco NDFC Release 12.1.2e, you can change persistent IP addresses which are assigned for mandatory pods such as POAP-SCP and SNMP trap. Refer Changing Persistent IP Address, on page 43.

This setting determines the Persistent IPs usage for the PODs required for Nexus Dashboard Fabric Controller. When you select Fabric Controller persona for the first time then, there is a precheck to see if Persistent IPs are allocated on Nexus Dashboard. If Persistent IPs are not allocated, then the operator sees an error. You can provide Persistent IPs in either Nexus Dashboard Management Network or Nexus Dashboard Data Network. Based, on this selection, you must specify the option under LAN Device Management Connectivity which can be found under Server Settings of NDFC application page. By default, Management is selected, but, if you provide Persistent IPs under Nexus Dashboard Data Network then, you must select Data as an option.



**Note** When you change the LAN Device Management connectivity from management to DATA or conversely. Some of the devices might have a CRITICAL Alert of 'SSH Unreachable' error for short time and eventually restored.

### SMTP Host under SMTP

This setting is used as EMAIL out-of-band notification for Programmable reports and Alarms. Starting NDFC 12.0.1a release, you can now receive NDFC Alarms and Reports through EMAIL notification. The SMTP Host address must be reachable through Nexus Dashboard Management Interface. If the Nexus Dashboard management interface and SMTP Host are part of different IP subnets then you must create a static route entry in Nexus Dashboard Cluster configuration.

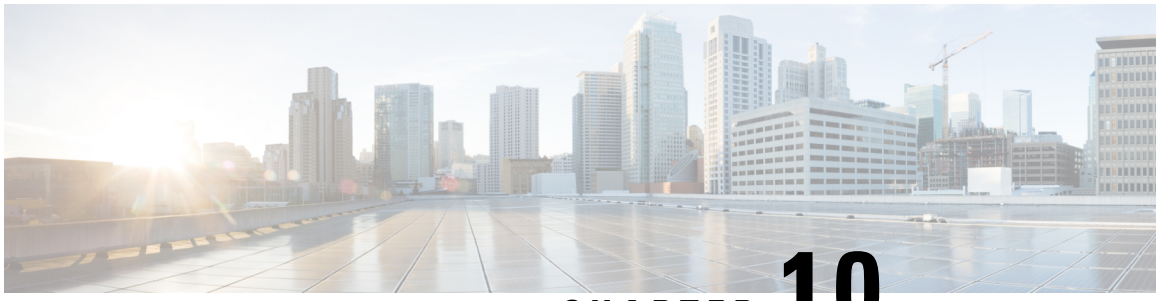
You can enter other texts for STMP fields. To initiate alarms to external receiver, provide IP of SNMP Listener and Port it is listening on.

### Disable Deployment Across all Fabrics Under LAN Fabric

This setting disables deployments for all the fabrics that are defined in the NDFC instance. You will not be able to enable the deployment on per fabric level. For example, if you have 3 fabrics then all 3 fabrics will be disabled from configuration point of view. You can continue to stage various configurations if necessary. Later, you can enable the deployment action by unchecking this server setting.

### Collect Temperature for LAN Switches Under PM

This setting enables to collect switch temperate details and show it in the Fabric Overview and then Metric section. By, default temperature data is not collected. Upon enabling this setting, you can view the temperature information of the fabric switches as well.



CHAPTER 10

# Feature Management


• Feature Management, on page 359


# Feature Management

In Cisco DCNM Release 11.x, you must choose the install mode while installing the DCNM. From Release 12.0.1a, Cisco Nexus Dashboard Fabric Controller allows you to install the service on the Nexus Dashboard. After you launch the Nexus Dashboard Fabric Controller UI, you will see three different Install modes on the Feature Management page.

Nexus Dashboard Fabric Controller 12 allows you to dynamically enable the feature set and scale applications. Choose **Settings > Feature Management** to choose the installer type and enable or disable few features on the selected deployment.

When you launch Nexus Dashboard Fabric Controller for the first time from Cisco Nexus Dashboard, the Feature Management screen appears. You can perform only Backup and Restore operations before you choose the feature set.

On the Feature Management page, you can choose one of the following install modes:

- Fabric Discovery
- Fabric Controller
- SAN Controller

After you select a Feature Set, from the next login, Dashboard page opens when you launch Cisco Nexus Dashboard Fabric Controller from Nexus Dashboard.

## Choosing Feature Set

When you launch Cisco Nexus Dashboard Fabric Controller 12 for the first time, none of the feature set is enabled. During this state, you can perform Backup and Restore to restore the DCNM 11.5(x) data on Nexus Dashboard Fabric Controller 12. Nexus Dashboard Fabric Controller will read the data from the backup file and select the installer type accordingly.

To deploy feature-set from Cisco Nexus Dashboard Fabric Controller Web UI perform the following steps:

**Procedure**

**Step 1** Choose **Settings** > **Feature Management**.

**Step 2** Select a persona to view the default set of features.

For information about the features available in Cisco NDFC personas, see Features with each Persona, on page 360.

**Step 3** In the table below, select the check box against the feature name available with the feature set.

**Step 4** Click **Apply**.

The feature-set will be deployed. The selected applications will be enabled. A message appears that the feature set is installed, and you must refresh to take effect.

**Step 5** Refresh the browser to deploy Nexus Dashboard Fabric Controller with the selected feature set and applications.

The left pane shows the features supported specifically with the deployed feature set.

## Features with each Persona

# Changing across Feature-Set

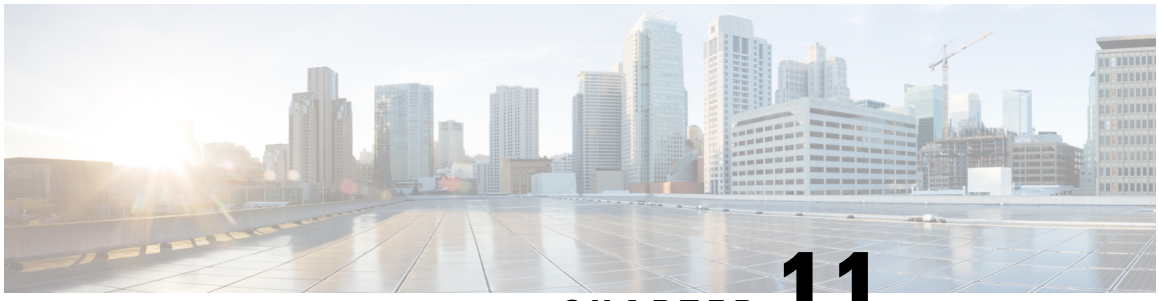
Nexus Dashboard Fabric Controller 12 allows you to switch from one feature set to another. Choose **Settings** > **Feature Management**. Select the desired feature set and applications in the table below. Click **Save & Continue**. Refresh the browser to begin using Cisco Nexus Dashboard Fabric Controller with the new feature set and applications.

There are a few features/applications supported with specific deployments. When you change the feature set, some of these features are not supported in the new deployment. The following table provides details about the pre-requisites and criteria based on which you can change the feature set.

***Table 40: Supported Switching between deployments***

| From/To | Fabric Discovery | Fabric Controller | SAN Controller |
|---|---|---|---|
| **Fabric Discovery** | - | Only monitor mode fabric is supported in Fabric Discovery deployment. When you change the feature set, the fabric can be used in the Fabric Controller deployment. | Not supported |
| **Fabric Controller** | You must delete the existing fabrics before changing the fabric set. | If you're changing from Easy Fabric to IPFM fabric application, you must delete the exiting fabrics. | Not supported |

| From/To | Fabric Discovery | Fabric Controller | SAN Controller |
|---|---|---|---|
| **SAN Controller** | Not supported | Not supported | - |



CHAPTER 11

# Credentials Management


- LAN Credentials Management, on page 363


# LAN Credentials Management

While changing the device configuration, Cisco Nexus Dashboard Fabric Controller uses the device credentials provided by you. However, if the LAN Switch credentials are not provided, Cisco Nexus Dashboard Fabric Controller prompts you to open the **Settings > LAN Credentials Management** page to configure LAN credentials.

Cisco Nexus Dashboard Fabric Controller uses two sets of credentials to connect to the LAN devices:

- **Discovery Credentials**—Cisco Nexus Dashboard Fabric Controller uses these credentials during discovery and periodic polling of the devices.

  NDFC used discovery credentials with SSH and SNMPv3 to discover hardware or software inventory from the switches. Therefore these are called as discovery credentials. You can discover one inventory per switch. These are read-only and cannot make configuration changes on the switches.

- **Configuration Change Credentials**—Cisco Nexus Dashboard Fabric Controller uses these credentials when user tries to use the features that change the device configuration.

LAN Credentials - You can use write option on LAN credentials to do configuration changes on the switch. One credential is allowed per user on a single switch. user-role must access to NDFC to use write option for the switches to push configuration on it through SSH connection.

For user-role created on NX-OS switches, an SNMPv3 user is created with same password. Ensure that the SSH and SNMPv3 credentials matches for the discovery of credentials. If SNMP authentication fails, discovery of credentials stops dislaying an error message. If SNMP authentication succeeds and SSH authentication fails, discovery of crendtials continues and the switch status displays a warning message for SSH error.

If user-role created on NX-OS switches uses AAA authentication, SNMPv3 user is not created. Using this AAA authentication to discover or import of a switch in NDFC the controller detects that the local SNMPv3 user is not created on the switch. Therefore, it runs exec command on the switch to create an SNMPv3 user with same password on the switch. The SNMPv3 user-role created is temporary. Once the user-role expires, continual discovery of switches from NDFC creates the SNMPv3 user.

LAN Credentials Management allows you to specify configuration change credentials. Before changing any LAN switch configuration, you must enter the LAN Credentials for the switch. If you do not provide the credentials, the configuration change action will be rejected.

These features get the device write credentials from LAN Credentials feature.

- Upgrade (ISSU)
- Maintenance Mode (GIR)
- Patch (SMU)
- Template Deployment
- POAP-Write erase reload, Rollback
- Interface Creation/Deletion/Configuration
- VLAN Creation/Deletion/Configuration
- VPC Wizard

You must specify the configuration change credentials irrespective of whether the devices were discovered initially or not. This is a one-time operation. After the credentials are set, the credentials will be used for any configuration change operation.

### Default Credentials

Default credentials is used to connect all the devices that the user has access to. You can override the default credentials by specifying credentials for each of the devices in the Devices below.

Cisco Nexus Dashboard Fabric Controller tries to use individual switch credentials in the Devices, to begin with. If the credentials (username/password) columns are empty in the Devices, the default credentials will be used.

### Switch Table

Devices table lists all the LAN switches that user has access. You can specify the switch credentials individually, that will override the default credentials. In most cases, you need to provide only the default credentials.

The LAN Credentials for the Nexus Dashboard Fabric Controller Devices table has the following fields.

| Field | Description |
|---|---|
| Device Name | Displays the switch name. |
| IP Address | Specifies the IP Address of the switch. |
| Credentials | Specifies whether default or switch specific custom credentials are used. |
| Username | Specifies the username that Nexus Dashboard Fabric Controller use to login. |
| Fabric | Displays the fabric to which the switch belongs. |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **LAN Credentials Management**.

| Action Item | Description |
|---|---|
| Edit | Choose a device name, click **Edit**, specify username and password. You can edit local or custom specific credentials |
| Clear | Choose a device name, click **Clear**.<br>A confirmation window appears, click **Yes** to clear the switch credentials from the NDFC server. |
| Validate | Choose a device name, click **Validate**.<br>A confirmation message appears, stating if the operation was successful or a failure. |

### Robot credentials

When you specify default credentials, you can enable the Robot feature. This enables the Robot flag.

Robot role is similar to earlier role in DCNM. The Robot user-role helps with switch and device accounting. You can track all the changes done on NDFC with a general user account. If the user-role changes on NDFC which impacts the change on the device which is termed as out-of-band changes. These changes are logged in the device as the changes made by a general user account. Therefore, you can track and distinguish between out-of-band changes and changes made on the device. This general user account is termed as robot user-role for the changes logged on the device.

For an example, a user-role with network-admin on NDFC has access to enter LAN device credential to push configuration on the switches. This user-role can check robot flag while creating LAN credentials.

The username mentioned for LAN credential is displayed on the changes logged in the device. If a username for LAN credential on NDFC is changed as controller and checks the robot flag, now the credentials for device changes from default to robot. This user-role pushes configuration on switches in NDFC. These changes are logged in history tab of fabric deployment as the changes made by user role network-admin, but the account logs on switch is showed as controller. Therefore, the appropriate user-role details are logged on NDFC and device.

In NDFC, robot user-role is considered as an admin role for all fabrics and devices. If default or credential is not set on a fabric you can use robot user-role, if it set for diferent devices. If other user-role with write access log into NDFC, this user-role will not be prompted to update the credentials as robot user-role is set. The credentials are set in order of an individual switch, robot and the default credentials

On **LAN Credentials Management** home page, you can choose either default credentials or robot credentials, while changing device configurations, unless customer credentials are set.

To set credentials, perform the following steps:

1. Choose required **Device Name** and click **Set**.

   The **Set Credentials** window appears.

2. Enter appropriate details. Choose **Robot** checkbox to set robot credentials.

   You can choose appropriate roles to push configurations to devices without adding device credentials

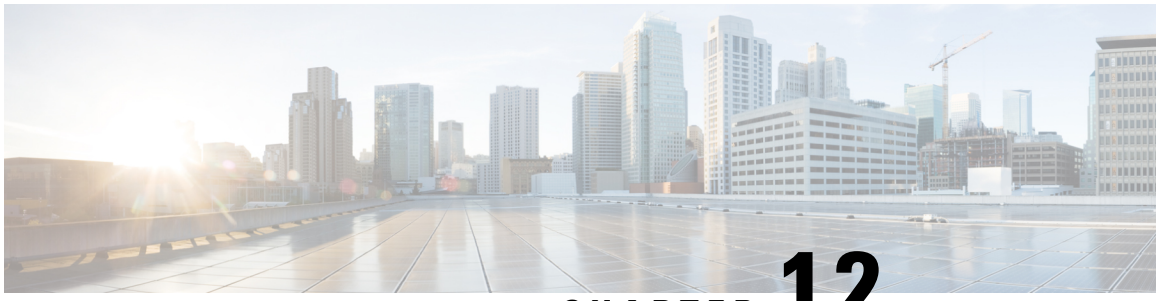
Choose required **Device Name** and click **Clear**. A confirmation message appears, click **Yes** to clear default device credentials.



# PART IV

# Operations


- Event Analytics, on page 369
- Image Management, on page 385
- Programmable Reports, on page 397
- License Management, on page 403
- Templates, on page 413
- Tech Support, on page 449
- Backup and Restore, on page 451
- NXAPI Certificates, on page 457



CHAPTER 12

# Event Analytics

This section contains the following topics:


- Alarms, on page 369
- Events, on page 378
- Accounting, on page 382
- Remote Clusters, on page 383


# Alarms

This tab displays the alarms that are generated for various categories. This tab displays information such as ID (optional), Severity, Failure Source, Name, Category, Acknowledged, Creation Time, Last Updated (optional), Policy, and Message. You can specify the Refresh Interval in this tab. You can select one or more alarms and then acknowledge or unacknowledge their status using the Change Status drop-down list. In addition, you can select one or more alarms and then click the Delete button to delete them.

## Alarms Raised

UI Path: **Operations > Event Analytics > Alarms**

Click the **Alarms Raised** tab to view the alarm policies that were triggered by an alarm.

Click on the required **Severity** column. A slide-in pane appears with policy severity details and description.

The following table describes the fields that appear on **Operations > Event Analytics > Alarms > Alarms Raised**.

| Field | Description |
|---|---|
| Severity | Specifies the severity of the alarm |
| Source | Specifies the name of the source. |
| Name | Specifies the name of the alarm |
| Category | Specifies the category of the alarm |
| Creation Time | Specifies the time at which the alarm was created |
| Policy | Specifies the policy of the alarm |
| Message | Displays the message. |

| Field | Description |
|---|---|
| Ack User | Displays the username who acknowledged the alarm. |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **Alarms Raised** tab.

| Action Item | Description |
|---|---|
| Acknowledge | Choose one or multiple alarms and choose **Acknowledge**. Allows you to bookmark the alarms and adds ack user name to **Acknowledged** column. |
| Unacknowledge | Choose one or multiple alarms and choose **Unacknowledge** to remove the bookmarked alarms.<br>**Note**<br>Only acknowledged alarms can be unacknowledged. |
| Clear | Choose alarm and choose **Clear** to clear the alarm policy manually.<br>The cleared alarms will be moved to **Alarm Cleared** tab. |
| Delete Alarm | Choose an alarm and choose **Delete** to delete the alarm. |



**Note** For link-down events, you must setup an external visible IP address for SNMP trap receiver, and configure switch to send SNMP trap to NDFC. Otherwise, the port state change can only be done through polling, which is every 5 minutes.

# Alarms Cleared

UI Path: **Operations > Event Analytics > Alarms > Alarms Cleared**

Alarms Cleared tab has the list of alarms which are cleared in the **Alarms Raised** tab. This tab displays information such as ID (optional), Severity, Failure Source, Name, Category, Acknowledged, Creation Time, Cleared At (optional), Cleared By, Policy, and Message. You can view the cleared alarm details for maximum of 90 days.

You can choose one or more alarms and click the **Actions > Delete** to delete them.

The following table describes the fields that appear on **Alarms Cleared** tab.

| Field | Description |
|---|---|
| Severity | Specifies the severity of the alarm. |
| Source | Specifies the IP Address of source alarm. |
| Name | Specifies the name of the alarm. |
| Category | Specifies the category of the alarm. |
| Creation Time | Specifies the time at which the alarm was created. |
| Cleared Time | Specifies the time at which the alarm was cleared. |

| Field | Description |
|---|---|
| Cleared By | Specifies the user who cleared the alarm. |
| Policy | Specifies the policy of the alarm. |
| Message | Specifies the CPU utilization and other details of alarm |
| Ack User | Specifies the acknowledged user role name. |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **Alarms Cleared** tab.

| Action Item | Description |
|---|---|
| Delete Alarm | Select an alarm and choose **Delete** to delete the cleared alarm |

# Monitoring and Adding Alarm Policies

In Cisco Nexus Dashboard Fabric Controller to enable alarms, Navigate to **Operations** > **Event Analytics** > **Alarms**, click **Alarm Policies** on vertical tab. Ensure that the Enable external alarms check box is selected. You must restart Nexus Dashboard Fabric Controller Server to bring this into effect.

You can forward alarms to registered SNMP Listeners in Nexus Dashboard Fabric Controller. From Cisco Nexus Dashboard Fabric Controller web UI, choose **Settings > Server Settings > Alarms**, ensure that the **Enable external alarms** check box is selected. You must restart Nexus Dashboard Fabric Controller Server to bring this into effect.

You can forward alarms to registered SNMP Listeners in Nexus Dashboard Fabric Controller. From Cisco Nexus Dashboard Fabric Controller web UI, choose **Settings > Server Settings > Alarms**, enter an external port address in alarm.trap.listener.address field, click **Apply Changes**, and restart SAN Controller.



**Note** Ensure that you select **Forwarding** check box in **Alarm Policy creation** dialog window to enable forwarding alarms to external SNMP Listener.

The following table describes the fields that appear on **Operations > Event Analytics > Alarms > Alarms Policies**.

| Field | Description |
|---|---|
| Name | Specifies the name of the alarm policy |
| Description | Specifies the description of the alarm policy |
| Status | Specifies the status of the alarm policy:<br>• Activated<br>• Deactivated |

| Field | Description |
|---|---|
| Policy type | Specifies the type of the policy:<br>• Device Health Policy<br>• Interface Health Policy<br>• Syslog Alarm Policy<br>• Hardware Health Policy |
| Devices | Specifies the devices to which the alarm policy is applied. |
| Interfaces | Specifies the interfaces. |
| Details | Specifies the details of the policy. |

The following table describes the action items, in the **Actions** menu drop-down list that appear on **Operations > Event Analytics > Alarms > Alarms Policies**.

| Action Item | Description |
|---|---|
| Create new alarm policy | Choose to create a new alarm policy. See Create new alarm policy section. |
| Edit | Select a policy and choose **Edit** to edit the alarm policy. |
| Delete | Select a policy and choose **Delete** to delete the alarm policy. |
| Activate | Select a policy and choose **Activate** to activate and apply the alarm policy. |
| Deactivate | Select a policy and choose **Deactivate** to disable and deactivate the alarm policy. |
| Import | Select to import alarm policies from a .txt file. |
| Export | • Click the box next to a specific alarm policy, then click **Export** to export that alarm policy as a .txt file.<br>• Select or deselect all the boxes next to the alarm policies, then click **Export** to export all the alarm policies as a .txt file. |

You can add alarm policies for the following:

- **Device Health Policy**: Device health policies enable you to create alarms when Device ICMP Unreachable, Device SNMP Unreachable, or Device SSH Unreachable. Also, these policies enable you to monitor chassis temperature, CPU, and memory usage.
- **Interface Health Policy**: Interface health policies enable you to monitor Up or Down, Packet Discard, Error, Bandwidth details of the interfaces. By default all interfaces are selected for monitoring.
- **Syslog Alarm Policy**: Syslog Alarm Policy defines a pair of syslog messages formats; one which raises the alarm, and one which clears the alarm.
- **Hardware Health Policy**: The hardware health policy is used to raise hardware-related alarms for different parameters, such as fan status, power supply, modular status and all interface-related alarms.

# Create new alarm policy

You can add alarm policies for the following:

- Device Health Policy
- Interface Health Policy
- Syslog Alarm Policy
- Hardware Health Policy

After you create a new alarm policy, in the **Alarm Policies** tab, click **Refresh** to view the newly-created alarm policy.

### Device Health Policy

Device health policies enable you to create alarms when certain conditions are met. By default, all devices are selected for monitoring.

- **Policy Name**: Specify a name for the policy. It must be unique.
- **Description**: Specify a brief description for the policy.
- **Forwarding**: You can forward alarms to registered SNMP listeners in Cisco Nexus Dashboard Fabric Controller . From the Web UI, choose **Settings** > **Server Settings** > **Events**.



**Note** Ensure that you select **Forwarding** check box while configuring alarm policies to forward alarms to an external SNMP listener.

- **Email**: You can forward alarm event emails to recipients when an alarm is created, cleared or when the severity is changed. From Cisco Nexus Dashboard Fabric Controller Web UI, choose **Settings** > **Server Settings** > **Events**. Configure the SMTP parameters, click **Save**, and restart Cisco Nexus Dashboard Fabric Controller services.
- Specify the CPU utilization parameters, memory utilization parameters, and environmental temperature parameters.
- **Device Availability**: Device health policies enable you to create alarms in the following situations:
  - **Device Access**: When device SNMP or device SSH is unnreachable.
  - **Peripherals**: When fan, power supply, or module is unnreachable.

Also, these policies enable you to monitor chassis temperature, CPU, and memory usage.

Select the devices for which you want to create policies. Specify the policy name, description, CPU Utilization parameters, Memory Utilization parameters, Environment Temperature parameters, device availability, and device features.

### Interface Health Policy

Interface health policies enable you to monitor the interface status, packet discards, errors and bandwidth details of the interfaces. By default, all interfaces are selected for monitoring.

Select the devices for which you want to create policies and then specify the following parameters:

- **Policy Name**: Specify a name for the policy. It must be unique.
- **Description**: Specify a brief description for the policy.
- **Forwarding**: You can forward alarms to registered SNMP listeners in Cisco Nexus Dashboard Fabric Controller . From the Web UI, choose **Settings** > **Server Settings** > **Events**.



**Note** Ensure that you select **Forwarding** check box while configuring alarm policies to forward alarms to an external SNMP listener.

- **Email**: You can forward alarm event emails to recipients when an alarm is created, cleared or when the severity is changed. From Cisco Nexus Dashboard Fabric Controller Web UI, choose **Settings** > **Server Settings** > **Events**. Configure the SMTP parameters, click **Save**, and restart Cisco Nexus Dashboard Fabric Controller services.
- **Linkstate**: Choose linkstate option to check for the interface link status. You can generate an alarm whenever a link is down and clear the alarms when the link is up.
- **Bandwidth (In/Out)**: Allows you to set the maximum bandwidth allowed in inbound and outbound directions. The system generates alarms when the bandwidth exceeds the specified values.
- **Inbound Errors**: Allows you to set thresholds for the number of inbound errors that are discarded after which it generates an alarm.
- **Outbound Errors**: Allows you to set thresholds for the number of outbound errors that are discarded after which it generates an alarm.
- **Inbound Discards**: Allows you to set thresholds for the number of inbound packets that are discarded after which it generates an alarm.
- **Outbound Discards**: Allows you to set thresholds for the number of outbound packets that are discarded after which it generates an alarm.

### Syslog Alarm

Syslog Alarm Policy defines a pair of Syslog messages formats; one which raises the alarm, and one which clears the alarm.

Select the devices for which you want to create policies and then specify the following parameters:

- Devices: Define the scope of this policy. Select individual devices or all devices to apply this policy.
- Policy Name: Specify the name for this policy. It must be unique.
- Description: Specify a brief description for this policy.
- Forwarding: You can forward alarms to registered SNMP listeners in Cisco Nexus Dashboard Fabric Controller . From Web UI, choose **Settings > Server Settings > Events**.



**Note** Ensure that you select **Forwarding** check box in Alarm Policy creation dialog window to enable forwarding alarms to external SNMP listener.

- Email: You can forward alarm event emails to recipient when alarm is created, cleared or severity changed. From Cisco Nexus Dashboard Fabric Controller Web UI, choose **Settings > Server Settings > Events**. Configure the SMTP parameters, click **Save**, and restart Cisco Nexus Dashboard Fabric Controller services.
- Severity: Define the severity level for this syslog alarm policy. Choices are: Critical, Major, Minor, and Warning.
- Identifier: Specify the identifier portions of the raise & clear messages.
- Raise Regex: Define the format of a syslog raise message. The syntax is as follows: Facility-Severity-Type: Message
- Clear Regex: Define the format of a syslog clear message. The syntax is as follows: Facility-Severity-Type: Message

The Regex definitions are simple expressions but not a complete regex. Variable regions of text are noted using $(LABEL) syntax. Each label represents a regex capture group (.+), which corresponds to one or more characters. The variable texts found in both raise and clear messages are used to associate the two messages. An Identifier is a sequence of one or more labels that appear in both messages. An Identifier is used to match a clear syslog message to the syslog message that raised the alarm. If the text appears only in one of the messages, it can be noted with a label and exclude it from the identifier.

**Example:** A policy with "Value": "ID1-ID2",

```
"syslogRaise": "SVC-5-DOWN: $(ID1) module $(ID2) is down $(REASON)"
"syslogClear": "SVC-5-UP: $(ID1) module $(ID2) is up."
```

In the example, ID1 and ID2 labels can be marked as an identifier to find the alarm. This identifier will be found in corresponding syslog messages. Label “REASON” is in the raise but not in the clear message. This label can be excluded from the identifier, as it has no impact on the syslog message to clear the alarm.

***Table 41: Example 1***

| **Identifier** | **ID1-ID2** |
|---|---|
| Raise Regex | ETHPORT-5-IF_ADMIN_UP: Interface Ethernet15/1 is admin up . |
| Clear Regex | ETHPORT-5-IF_DOWN_NONE: Interface Ethernet15/1 is down (Transceiver Absent) |

In the above example, the regex expressions are part of the syslog messages that appear in the terminal monitor.

***Table 42: Example 2***

| **Identifier** | **ID1-ID2** |
|---|---|
| Raise Regex | ETH_PORT_CHANNEL-5-PORT_DOWN: $(ID1): $(ID2) is down |
| Clear Regex | ETH_PORT_CHANNEL-5-PORT_UP: $(ID1): $(ID2) is up |

***Table 43: Example 3:***

| **Identifier** | **ID1-ID2** |
|---|---|
| Raise Regex | ETHPORT-5-IF_SFP_WARNING: Interface $(ID1), High Rx Power Warning |

| Identifier | ID1-ID2 |
|---|---|
| Clear Regex | ETHPORT-5-IF_SFP_WARNING: Interface $(ID1), High Rx Power Warning cleared |

### Hardware Health Policy

The hardware health policy is used to raise hardware-related alarms for different parameters, such as fan status, power supply, modular status and all interface-related alarms.

By default, there is a hardware policy called as discovery that is standard with the NDFC installation. This hardware policy defines various conditions for different parameters. You can also create custom hardware policies for the parameters listed above and define regex expressions based on which alarms are raised.

By default, the **All Devices** option is selected automatically.

- **Policy Name**: Specify a name for the policy. It must be unique.
- **Description**: Specify a brief description for the policy.
- **Forwarding**: You can forward alarms to registered SNMP listeners in Cisco Nexus Dashboard Fabric Controller . From the Web UI, choose **Settings** > **Server Settings** > **Events**.



**Note** Ensure that you select **Forwarding** check box while configuring alarm policies to forward alarms to an external SNMP listener.

- **Email**: You can forward alarm event emails to recipients when an alarm is created, cleared or when the severity is changed. From Cisco Nexus Dashboard Fabric Controller Web UI, choose **Settings** > **Server Settings** > **Events**. Configure the SMTP parameters, click **Save**, and restart Cisco Nexus Dashboard Fabric Controller services.

Hardware alarms are raised based on regex expressions that you enter when you are creating the policy.

In the **Alarms** area, create a hardware health policy to raise alarms for the following parameters:

- **Fan**: Define the severity for fan-related alarms and determine the condition for the alarms.
  1. Click the toggle switch next to **Fan** to enable the fan-related alarms.
  2. Select the severity of the alarm:
     - Critical
     - Major
     - Minor
     - Warning
     - Cleared
  3. Click **Edit Regex**, then enter the value that will trigger the alarm when the status is not that value.

     For example, if you enter `ok` in the **Trigger alarm when status is not** field, NDFC will raise an alarm for any status other than `ok`, such as N/A.

4. Click **Save**.

- **Power Supply**: Define the severity for power supply-related alarms and determine the condition for the alarms.
  1. Click the toggle switch next to **Power Supply** to enable the power supply-related alarms.
  2. Select the severity of the alarm:
     - Critical
     - Major
     - Minor
     - Warning
     - Cleared
  3. Click **Edit Regex**, then enter the value that will trigger the alarm when the status is not that value.

     For example, if you enter `ok` in the **Trigger alarm when status is not** field, NDFC will raise an alarm for any status other than `ok`, such as failed, OffEnvpower, OffDenied, and so on.
  4. Click **Save**.
- **Module**: Define the severity for module-related alarms and determine the condition for the alarms.
  1. Click the toggle switch next to **Module** to enable the module-related alarms.
  2. Select the severity of the alarm:
     - Critical
     - Major
     - Minor
     - Warning
     - Cleared
  3. Click **Edit Regex**, then enter the value that will trigger the alarm when the status matches that value.

     For example, if you were to enter the following value in the **Trigger alarm when status matches regex** field, as shown in the information (i) button:

     ```
     ^(?!ok|poweredDown|okButDiagFailed).*$
     ```

     NDFC will raise an alarm when modules are in states other than ok, poweredDown, and OkButDiag failed.
  4. Click **Save**.
- **Interface Status**: Define the severity for interface-related alarms and determine the condition for the alarms.
  1. Click the toggle switch next to **Interface Status** to enable the interface-related alarms.
  2. Click one or more toggle switches next to the appropriate severity to select the severity of the alarm:

- Critical
- Major
- Minor
- Warning
- Cleared

3. Click **Edit Regex**, then enter the value that will trigger the alarm when the status matches that value. The provided regex expression is matched against the combined field of `admin_status:oper_status:status_reason`.

   For example, if you were to enter the following value in the **Trigger alarm when status matches regex** field:

   ```
   ^up:down:(?!Link not connected|XCVR not inserted|sfpNotPresent|Channel admin down).*$
   ```

   NDFC will raise an alarm when interfaces are in states that match these values.

4. Click **Save**.

# Events

This tab displays the events that are generated for the switches. This tab displays information such as Ack, Acknowledged user, Group, Switch, Severity, Facility, Type, Count, Last Seen, and Description. You can select one or more events and then acknowledge or unacknowledge their status using the Change Status drop-down list. In addition, you can select one or more alarms and then click the Delete button to delete them. If you want to delete all events, click the Delete All button.

The following table describes the fields that appear on **Operations > Event Analytics > Events**.

| Field | Description |
|---|---|
| Group | Specifies the Fabric |
| Switch | Specifies the hostname of the switch |
| Severity | Specifies the severity of the event |
| Facility | Specifies the process that creates the events.<br>The event facility includes two categories: NDFC and syslog facility. Nexus Dashboard Fabric Controller facility represents events generated by Nexus Dashboard Fabric Controller internal services and SNMP traps generated by switches. Syslog facility represents the machine process that created the syslog messages. |
| Type | Specifies how the switch/fabric are managed |
| Count | Specifies the number of times the event has occurred |
| Creation Time | Specifies the time when the event was created |
| Last Seen | Specifies the time when the event was run last |

| Field | Description |
|---|---|
| Description | Specifies the description provided for the event |
| Ack | Specifies if the event is acknowledged or not |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **Operations > Event Analytics > Events**.

| Action Item | Description |
|---|---|
| Acknowledge | Select one or more events from the table and choose **Acknowledge** icon to acknowledge the event information for the fabric.<br>After you acknowledge the event for a fabric, the acknowledge icon is displayed in the Ack column next to the Group. |
| Unacknowledge | Select one or more events from the table and choose **Unacknowledge** icon to acknowledge the event information for the fabric. |
| Delete | Select an event and choose **Delete** to delete the event. |
| Add Suppressor | Select an event and choose **Add Suppressor** to add a rule to the event. You can provide name to the rule. Using the **Scope** options, you can add this rule to all the Fabrics, or particular elements or all elements. |
| Event Setup | Allows you to setup new event. For more information, see Event Setup, on page 379. |

# Event Setup

To setup an event using the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **Operations** > **Event Analytics** and click on the **Events** tab.

**Step 2** From the **Actions** drop-down list, select **Event Setup**.
The **Receiver** tab displays the following details:

- **Syslog Receiver enabled**: Displays the status of the syslog server.
- **SNMP Trap Receiver**: Displays the details of SNMP traps received, processed and dropped.
- **Syslog Receiver**: Displays the details of syslog messages received, processed and dropped.

**Step 3** Navigate to the **Sources** tab, to view a list of fabrics and its associated switches.
The **Sources** tab displays all the fabrics and the associated switches in tabular format. It also displays if traps and syslogs have been configured on the switches.

**Step 4** Perform the following steps to create rules for forwarding email notifications or traps for events:

Cisco Nexus Dashboard Fabric Controller Web UI forwards fabric events through email or SNMPv1 or SNMPv2c traps. Some SMTP servers may require adding authentication parameters to the emails that are sent from Nexus Dashboard Fabric Controller to the SMTP servers.

a) Ensure that you have configured SMTP parameters before configuring rules for forwarding event notifications through emails. To verify SMTP configuration, navigate to **Settings** > **Server Settings** > **SMTP** and verify that you have configured the required fields.

a) To enable events forwarding, choose **Settings** > **Server Settings** > **Events** and configure the fields as described in the following table.

*Table 44: Configure Events Forwarding*

| Field | Description |
|---|---|
| **Enable Event forwarding** | Check the checkbox to enable events forwarding feature. |
| **Email Forwarding From Email List** | Specifies the email address from which the forwarding messages arrive. |
| **Snooze Event Forwarding** | Snoozes an event from forwarding for the given time range. |
| **Maximum Number of Repeats in Event Forwarding** | Stops forwarding an event after the specified time. 0 indicates unlimited time. |
| **Maximum Number in Events/Traps/Syslog Queue** | Specifies the maximum number in the queue before dropping the incoming events/traps/syslog. |

b) To configure rules, choose **Operations** > **Event Analytics**.

c) Navigate to the **Forwarding** tab and choose **Actions** > **Add Rule** and configure the fields as described in the following table.

*Table 45: Configure Rules*

| Field | Description |
|---|---|
| **Forwarding Method** | Chooose one of the forwarding methods:<br>• **E-Mail**<br>• **Trap** |
| **Email Address** | This field appears if you select **E-mail** as the forwarding method.<br>Enter an email address for forwarding the event notifications. |
| **Address** | This field appears if you select **Trap** as the forwarding method.<br>Enter the IP address of the SNMP trap receiver. You can either enter an IPv4 or IPv6 address or a DNS server name. |

| Field | Description |
|---|---|
| **Port** | Enter the port to which the traps are forwarded. |
| **Forwarding Scope** | Maximum number in queue before dropping the incoming events/traps/syslog messages. |
| **Fabric** | Select **All Fabrics** or a specific fabric for notification. |
| **Source** | Select **DCNM** or **Syslog**.<br>If you select **DCNM**, do the following:<br>1. From the **Type** drop-down list, choose an event type.<br>2. Check the **Storage Ports Only** check box to select only the storage ports. This check box is enabled only for port related events.<br>If you select **Syslog**, do the following:<br>1. In the **Facility** list, select the syslog facility.<br>2. In the **Type** field, enter the syslog type.<br>3. In the **Description Regex** field, enter a description that matches with the event description. |

d) From the **Minimum Severity** drop-down list, select the severity level of the messages to receive.

The traps that are transmitted by Cisco Nexus Dashboard Fabric Controller correspond to the severity type. A text description is also provided with the severity type.

```
trap type(s) = 40990 (emergency)
40991 (alert)
40992 (critical)
40993 (error)
40994 (warning)
40995 (notice)
40996 (info)
40997 (debug)
textDescriptionOid = 1, 3, 6, 1, 4, 1, 9, 9, 40999, 1, 1, 3, 0
```

e) Click **Add Rule**.

**Step 5** Perform the following steps to create rules for suppressing events:

Nexus Dashboard Fabric Controller allows you to suppress specified events based on user-specified rules. Such events will not be displayed on the Nexus Dashboard Fabric Controller Web UI and SAN Client. The events will neither be added to the Nexus Dashboard Fabric Controller database, nor forwarded via email or as SNMP traps.

You can view, add, modify, and delete rules from the table. You can create a rule from the existing events. Select an existing event as the template and open the **Add Rule** window by navigating to **Operations** > **Event Analytics** > **Events** page, select the event and choose **Actions** > **Add Suppresor**. The details are automatically ported from the selected event in the events table to the fields of the **Add Rule** window.

a) In the **Name** field, enter a name for the rule.
b) In the **Scope** field, select one of the following options - **SAN**, **Port Groups** or **Any**.

   In the **Scope** field, the LAN/SAN groups and the port groups are listed separately. For SAN and LAN, select the scope of the event at the fabric or group or switch level. You can only select groups for port group scope. If use select **Any** as the scope, the suppression rule is applied globally.

c) In the **Facility** field, enter the name or choose from the SAN/LAN switch event facility list.

   If you do not specify a facility, a wildcard is applied.

d) In the **Type** field, enter the event type.

   If you do not specify the event type, wildcard is applied.

e) In the **Description Matching** field, specify a matching string or regular expression.

   The rule matching engine uses regular expression that is supported by Java Pattern class to find a match against an event description text.

f) Check the **Active Between** check box and select a valid time range during which the event is suppressed.

   By default, the time range is not enabled.

   **Note**
   In general, you must not suppress accounting events. Suppression rule for Accounting events can be created only for certain situations where accounting events are generated by actions of Nexus Dashboard Fabric Controller or switch software. For example, 'sync-snmp-password' AAA syslog events are automatically generated during the password synchronization between Nexus Dashboard Fabric Controller and managed switches. To suppress accounting events, navigate to **Operations** > **Event Analytics** > **Events** page, select the event and choose **Actions** > **Add Suppressor**.

g) Click **Add Rule**.

# Accounting

You can view the accounting information on Cisco Nexus Dashboard Fabric Controller Web UI.

The following table describes the fields that appear on **Operations > Event Analytics > Accounting**.

| Field | Description |
|---|---|
| Source | Specifies the source |
| User Name | Specifies the user name. |
| Time | Specifies the time when the event was created |
| Description | Displays the description. |
| Group | Specifies the name of the group. |

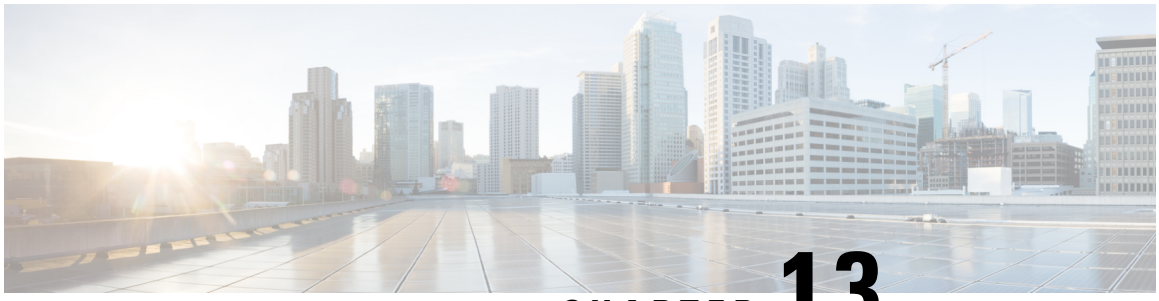
The following table describes the action items, in the **Actions** menu drop-down list, that appear on **Operations > Event Analytics > Accounting**.

| Action Item | Description |
| --- | --- |
| Delete | Select a row and choose **Delete** to delete accounting information from the list. |

# Remote Clusters

This tab displays the clusters and the number of Fabrics in each cluster in your setup.

Click on the Cluster Name to see the summary information. You can click on the launch icon to view the detailed summary of the Cluster.



CHAPTER 13

# Image Management


- Image Management, on page 385


# 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.
- In order to execute any ISSU operations, any new NDFC user must first set the necessary device credentials under the Credential Management page. You will not be able to execute ISSU operations without first setting the proper device credentials.

The **Image Management** window has the following tabs and you can perform the operations listed in the Actions column.

| Tabs | Actions |
|---|---|
| Overview | You can view dashlets for uploaded image and related information. |
| Images | Upload |
| Image Policies | Create |
| History | History, on page 396 |

Ensure that your user role is **network-admin** or **device-upg-admin** and you didn't freeze the Nexus Dashboard Fabric Controller 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 Nexus Dashboard Fabric Controller is in freeze mode.

Here's the process to upgrade the switch image:

1. Discover the switches into Nexus Dashboard Fabric Controller.
2. Upload images.
3. Create image policies.
4. Attach the image policies to the switches.
5. Stage the images on switches.
6. (Optional) Validate if the switches support non-disruptive upgrade.
7. Upgrade the switches accordingly.

# Overview

Image Management deploys Cisco software images to switches which allows network stability and feature consistency. The benefits of Image Management workflows include the following:

- Comprehensive image management policies allow you to specify a version and patch level for a set of switches
- Nexus Dashboard Fabric Controller (NDFC) validates compliance for the image policy associated with each switch
- Image staging, validation, and in-service software upgrade (ISSU) operations are independent, allowing mass upgrades and downgrades



**Note** Starting with Nexus Dashboard Fabric Controller release 12.1.3, NDFC provides the ability to perform staging and validation in a single step.

- You can perform the following operations before the maintenance window:
  - Staging image files

    This copies the image files to the bootflash.
  - Validate Network Operating System (NOS) and EPLD compatibility where possible

    This checks if the image is complete, if the image is valid for the individual hardware, and if the upgrade can be non-disruptive.

        - Run pre-upgrade report

- The ability to run reports pre/post-upgrade and compare the results
- The ability to generate snapshots of the configuration pre/post-upgrade
- The **View Details** column provides Live log status to monitor each operation
- Allow users to make use of maintenance mode to minimize the impact of disruptive upgrades, especially for multi-reload upgrade situations
- Upgrade groups allows bulk upgrades and downgrades. Upgrade groups have checks to avoid unnecessary downtime in redundant fabrics in the following cases:
    - All switches of a given role for a fabric will be in the same group
    - All route reflectors (RRs) in a fabric will be in the same group
    - All rendezvous points (RPs) in a fabric will be in the same group
    - Both Virtual Port Channel (vPC) peers will be in the same group
    - All In-Band Seed Devices will be in the same group
- Give visibility into previous and current upgrade details as well as high level summarization
- Visibility into current NOS, EPLD and patch consistency at a switch, fabric, and group level

Starting with Nexus Dashboard Fabric Controller release 12.1.2, the **Overview** tab has changed and displays the images, policies, fabric status, and switch upgrade group status.

The Image Management UI has the following functional areas:

- **Overview**: This displays the images, policies, fabric status, and switch upgrade group status.
    - The **Images** card displays the number of images and the type of packages or patches.
    - The **Policies** card displays the number of polices, platforms and release versions.
    - The **Fabric Status** card displays the number of devices in the fabric and if the devices are in or out of sync and displaying in either red or green.
    - The **Switch Upgrade Group Status** card displays the number of groups and device status.
        - Green: Devices that are in sync.
        - Red: Devices that are out of sync.
        - Gray: Devices that are not available (N/A). An example is when a device does not have a policy set.
- **Images**: This displays the images. You can upload or delete images.
- **Image Polices**: This displays the image policies. You can create, delete, or edit image policies.
- **Devices**: This displays the devices. You can stage image, upgrade, validate, change mode, attach group, detach group, attach policy, or detach policy.
- **History**: This displays the history of all the operations performed on the switches.

# Images

You can view the details of the images and the platform under this tab. You can upload or delete images to a device.

The following table describes the fields that appear on **Operations > Image Management > Images**.

| Field | Description |
|---|---|
| Platform | Specifies the name of the platform. Images, RPMs, or SMUs are categorized as follows:<br>• N9K/N3k<br>• N6K<br>• N7K<br>• N77K<br>• N5K<br>• Other<br>• Third Party<br>The images are the same for N9K and N3K platforms.<br>The platform is **Other** if the uploaded images are not mapped to any of the existing platforms.<br>The platform is **N9K/N3K** for RPMs. |
| Bits | Specifies the bits of the image |
| 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 Sub Type | Specifies the file type of the image, EPLD, RPM, or SMU.<br>The file type EPLDs are **epld**. The file types of images are **nxos**, **system** or **kickstart**. The file type for RPMs is **feature** and for SMUs the file type is **patch**. |
| NXOS Version | Specifies the NXOS image version for only Cisco switches. |
| Image Version | Specifies the image version for all devices, including the non-Cisco devices as well. |
| Size (Bytes) | Specifies the size of the image, RPM, or SMU files in bytes. |

| Field | Description |
|---|---|
| Checksum | Specifies the checksum of the image. The checksum checks if there's any corruption in the file of the image, RPM, or SMU. You can validate the authenticity by verifying if the checksum value is same for the file you downloaded from the Cisco website and the file you upload in the **Image Upload** window. |

The following table describes the action items, in the **Actions** menu drop-down list, that appears on **Operations > Image Management > Images**.

| Action Item | Description |
|---|---|
| Refresh | Refreshes the Images table. |
| Upload | Click to upload a new image. For instructions, see Uploading an Image, on page 389. |
| Delete | Allows you to delete the image from the repository.<br>Choose an image, click **Actions**, and choose **Delete**. A confirmation window appears. Click **Yes** to delete the image.<br>**Note**<br>Before deleting an image, ensure that the policy attached to the image, is not attached to any switches.<br>**Note**<br>If you delete an image on a switch in switch console, allow maximum of 24 hours to refresh and view update on NDFC. Else, on NDFC UI, navigate **LAN** > **Fabrics** > **Switches**, choose switch for which image is deleted and click **Actions** > **Discover** > **Rediscover** to view updates. |

## Uploading an Image

You can upload 32-bit and 64-bit images. To upload different types of images to the server from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Note** Devices use these images during POAP or image upgrade. All the images, RPMs, and SMUs are used in the **Image 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 **Operations > Image Management > Images**.

**Step 2** Click **Actions** and choose **Upload**.

The **Upload Image** 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 **OK**.

You can upload a ZIP or TAR file as well. Cisco Nexus Dashboard Fabric Controller 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>`.

**Note**
If only EPLD files are uploaded, you cannot create policy as Release drop-down list is empty for EPLD images.

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 Nexus Dashboard Fabric Controller to 1 hour so that the image upload is complete. To increase the wait time from Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

a) Choose **Settings > Server Settings**.
b) Search for the **csrf.refresh.time** property, and set the value as **60**.

The value is in minutes.

c) Click **Apply Changes**.
d) Restart the Nexus Dashboard Fabric Controller server.

# Image 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. Based on the policy applied on a switch, Cisco Nexus Dashboard Fabric Controller 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 describes the action items, in the **Actions** menu drop-down list, that appear on **Operations > Image Management > Image Policies**.

| Action Item | Description |
|---|---|
| Create | Allows you to create a policy that can be applied to images. See Creating an Image Policy, on page 391 section. |
| Delete | Allows you to delete the policy.<br>Choose a policy, click **Actions**, and choose **Delete**. A confirmation window appears. Click **Confirm** to delete the **policy**.<br>**Note**<br>An error message appears if you try to delete a policy that is attached to a device. |
| Edit | Allows you to edit the policy. |

## Creating an Image Policy

To create an image policy from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Before you begin

Upload the images under the **Images** tab before creating an image policy. See the Uploading an Image, on page 389 for more information about uploading images.

### Procedure

**Step 1** Choose **Operations** > **Image Management** > **Image Policies**.

**Step 2** Click **Actions** > **Create**.

The **Create Image Management Policy** dialog box appears.

**Step 3** Enter information for the required fields.

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 **Images** window. The options for the **Release** drop-down list will be autopopulated based on the platform you choose. |

| Fields | Actions |
|---|---|
| Release | Choose the NX-OS version from the **Release** drop-down list.<br>The release versions of 64-bit images are appended with 64bit in the image name.<br>**Note**<br>If only EPLD files are uploaded, you cannot create policy as Release drop-down list is empty for EPLD images. |
| Package Name | (Optional) Choose the packages. before choose Packages, View All Packages check box to display all uploaded packages for a given platform (its version agnostic). |
| Policy Description | (Optional) Enter a policy description. |
| EPLD | (Optional) Check the EPLD check box if the policy is for an EPLD image. |
| Select EPLD | (Optional) Choose the EPLD image. |
| RPM Disable | (Optional) Check this check box to uninstall the packages. |
| RPMs To Be Uninstalled | (Optional) Enter the packages to be uninstalled separated by commas. You can enter the package names only if you check the RPM Disable checkbox. |

**Step 4** Click **Save**.

### What to do next

- Attach the policy to a device. See Modifying a Policy, on page 395 section for more information.
- To edit an image policy after you've created it, click **Actions** > **Edit**.
- To delete an image policy, click **Actions** > **Delete**.

# Devices

The **Devices** window displays all the switches that you discover in the Cisco Nexus Dashboard Fabric Controller. You can view information like the current version of the switch, policy attached to it, status, and other image-related information. You can filter and sort the entries.

You can click on **Policy** column to view associated policy information for the switch. Similarly, click on **View details** column for required switch to view details. The view details columns can have either **Validate**, or **Upgrade**, or **None**.

You can perform the following actions in the **Devices** window:

## Staging an Image

After attaching an image policy to a switch, stage the image. When you stage an image, the files are copied into the bootflash.

To stage an image from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Before you begin

- Attach a policy to the selected devices before staging an image on the device.
- The minimum supported NX-OS image version in Fabric Controller is 7.0(3)I7(9).

  To stage an image on Cisco Nexus 9000 or Nexus 3000 switches running NX-OS version earlier than the version mentioned above, you must set **Use KSTACK to SCP on N9K, N3K** value to False. On the Web UI, choose **Settings > Server Settings > SSH** tab. Uncheck the **Use KSTACK to SCP on N9K, N3K** check box. If you're staging supported image versions, check this check box.

### Procedure

**Step 1** Choose **Operations > Image Management > Devices**.

**Step 2** Choose a switch by checking the check box.

**Note**
You can choose more than one switch to stage an image.

**Step 3** Click **Actions** and choose **Stage Image**.

The **Select Images to Install window** appears.

In this window, you can view how much space is available on the switch and how much space is required.

**Step 4** (Optional) Click the hyperlink under the Files For Staging column to view the files that are getting copied to the bootflash.

**Step 5** Click **Stage**.

You will be diverted to the Devices tab under the Image Management window.

**Step 6** (Optional) You can view the status under the Image Staged column.

**Step 7** (Optional) Click the hyperlink under the Reason column to view the log.

## Validating an Image

Before you upgrade the switches, you can validate if they support non-disruptive upgrade. To validate an image from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Procedure

**Step 1** Choose **Operations > Image Management > Devices**.

**Step 2** Choose a switch by checking the check box.

**Note**
You can choose more than one switch to stage an image.

**Step 3** Click **Actions** and choose **Validate**.

The **Validate** dialog box appears.

**Step 4** Check the Confirm non disruptive upgrade check box.

**Step 5** Click **Validate**.

You'll return to the Overview tab under the Image Management window.

**Step 6** (Optional) You can view the status under the Validated column.

**Step 7** (Optional) Click the hyperlink under the Reason column to view the log.

## Upgrading an Image

You can upgrade or uninstall a switch. Upgrade Groups option allows you to trigger image upgrade on multiple switches at an instant. This option can be selected for upgrade/downgrade options.



**Note** It is recommended to perform upgrade for maximum of twelve switches at once. If you choose more than twelve switches, the upgrade happens sequentially.

**Upgrade Options for NX-OS Switches**

- Disruptive: Choose this option for disruptive upgrades.
- Allow Non-disruptive: Choose this option to allow non-disruptive upgrades. When you choose **Allow Non Disruptive** option and if the switch does not support non-disruptive upgrade, then it will go through a disruptive upgrade. When you choose **Force Non Disruptive** and 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.

To upgrade a switch image from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Procedure

**Step 1** Choose **Operations** > **Image Management** > **Devices**.

**Step 2** Choose a switch by checking the check box.

**Step 3** Click **Actions** and choose **Upgrade**.

The **Upgrade/Uninstall** window appears.

**Step 4** Choose the type of upgrade by checking the check box.

The valid options are NXOS, EPLD, and Packages (RPM/SMU).

**Step 5** Choose NXOS, EPLD, or Packages:

a) Choose an upgrade option from the drop-down list based on how you want to upgrade.

b) (Optional) Check the BIOS Force check box.

You can view thee validation status of all the devices.

c) Check the **Golden** check box to perform a golden upgrade.
d) Enter the module number in the **Module Number** field.

You can view the module status below this field.

**Note**
- If you choose **Packages**, you can view the package details too.
- You can uninstall the packages by selecting the **Uninstall** radio button.

**Step 6** Click **Upgrade**.

**Note**
Upgrade status takes 30 - 40 minutes to update, if multiple switches are upgraded.

## Modifying a Policy

You can update the image policy that you have attached to a switch. You can change an image policy for multiple switches at the same time.

To attach or change an image policy attached to a switch from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **Operations** > **Image Management** > **Devices**.
**Step 2** Choose a switch by checking the check box.
**Step 3** Click **Actions** and choose **Modify Policy**.

The **Apply Policy** dialog box appears.

**Step 4** You can either attach or detach a policy, choose required check box.
**Step 5** Choose a policy from the Policy drop-down list.
**Step 6** Click required **Attach** or **Detach**.
**Step 7** (Optional) Click the hyperlink under the Reason column to view the changes.
**Step 8** (Optional) Click the hyperlink under thee Status column to view the current and expected image versions.

If the switch is in **Out-Of-Sync** status, view the expected image versions and upgrade the switch accordingly.

## Recalculating Compliance

To recalculate the configuration compliance of a switch from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **Operations** > **Image Management** > **Devices**.

**Step 2** Choose a switch by checking the check box.

**Step 3** Click **Actions** and choose **Recalculate Compliance**.

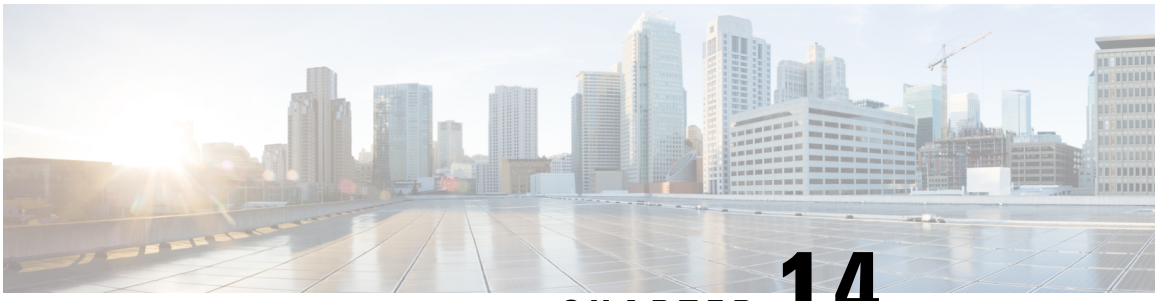
**Step 4** Click the hyperlink under the Reason column to view the changes.

# History

You can view the history of all the Image Management operations from **Operations** > **Image Management** > **History** tab.

The following table describes the fields that appear on this screen.

| Field | Description |
|---|---|
| ID | Specifies the ID number. |
| Device Name | Specifies the device name. |
| Version | Specifies the version of the image on the device. |
| Policy Name | Specifies the policy name attached to the image. |
| Status | Displays if the operation was a success or failure. |
| Reason | Specifies the reason for the operation to fail. |
| Operation Type | Specifies the type of operation performed. |
| Fabric Name | Specifies the name of the Fabric. |
| Created By | Specifies the user name who performed the operation. |
| Timestamp | Specifies the time when the operation was performed. |



CHAPTER 14

# Programmable Reports

The **Programmable Reports** application enables the generation of reports using Python 2.7 scripts. Report jobs are run to generate reports. Each report job can generate multiple reports. You can schedule the report to run for a specific device or fabric. These reports are analyzed to obtain detailed information about the devices.

The **REPORT** template type is used to support the **Programmable Reports** feature. This template has two template subtypes, **UPGRADE** and **GENERIC**. For more information on the REPORT template, refer Report Template, on page 440. A python SDK is provided to simplify report generation. This SDK is bundled with Nexus Dashboard Fabric Controller.

**Note** A Jython template supports a maximum file size of 100k bytes. In case any report template exceeds this size, Jython execution may fail.

### Nexus Dashboard Fabric Controller UI Navigation

To launch programmable reports on the Cisco Nexus Dashboard Fabric Controller Web UI, choose **Operations** > **Programmable Reports**.

The **Reports** window is displayed. This window has **Report Definitions** and **Reports** tabs. You can create reports from both the tabs by clicking **Create Report**. For information on creating a report job, refer *Creating a Report Job*. Refresh the window by clicking the **Refresh** icon.

**Note** Report jobs and SAN user defined reports are not migrated when upgraded from Cisco DCNM 11.5(x) to Nexus Dashboard Fabric Controller Release 12.0.1a. You must create them again manually.

This chapter contains the following sections:


- Create Report, on page 398
- Report Definitions, on page 399
- Reports, on page 401

# Create Report

Choose **Operations** > **Programmable Reports**. Click **Create Report**. The **Create Report** wizard appears.

To create a report job, perform the following steps:

### Procedure

**Step 1** Enter a name for the report job in the **Report Name** field.

**Step 2** Click **Select a template**.

**Step 3** Choose a report template from the drop-down list and click **Select**.

Based on the template you've chosen, provide required values to the fields that appear on the screen.

**Step 4** Click **Next** to move to the **Source & Recurrence** step.

**Step 5** Choose the frequency at which the report job should be run.

The following table shows the options available and their description.

| Available Option | Description |
|---|---|
| Now | The report is generated now. |
| Daily | The report is generated daily at a specified time between the Start Date and End Date. |
| Weekly | The report is generated once a week at a specified time between the Start Date and End Date. |
| Monthly | The report is generated once a month at a specified time between the Start Date and End Date. |
| Periodic | The report is generated periodically in a time period between the specified Start Date and End Date. The interval of time between the reports can be specified in minutes or hours. |

**Step 6** In the **Email Report To** field, enter an email ID or mailer ID if you want the report in an email.

You must configure SMTP settings in **Settings > Server Settings > SMTP** tab. If the Data service IP address is in private subnet, the static management route for SMTP server must be added in Cisco Nexus Dashboard cluster configuration.

**Step 7** Choose the devices, fabrics, or VSANs in the **Select device(s)**, **Select fabric(s)**, or **Select VSAN(s)** area.

**Note**
Based on the template you choose, the devices, fabrics, or VSANs are populated.

**Step 8** Click **Save**.

A new report and repot definitions are created and appears on the **Reports** and **Report Definitions** tab respectively.

# Report Templates

Each report template has some data associated with it. Depending on the features you have enabled in Nexus Dashboard Fabric Controller, some of the report templates available are

- Inventory_Report
- Performance_Report
- Switch_Performance_Report
- fabric_cloudsec_oper_status
- fabric_macsec_oper_status
- fabric_nve_vni_counter
- fabric_resources
- sfp_report
- switch_inventory

In addition to the templates listed above, any other templates that have been created by you will also be listed here. For more information on default templates and creating customized templates, refer to Templates, on page 413. Templates are listed based on the associated tags.

**Performance_Report** and **Switch_Performance_Report** are used for performance management reports.

# Report Definitions

The **Report Definitions** tab displays the report jobs which are created by a user.

You can view the following information in this tab:

| Field | Description |
|---|---|
| Title | Specifies the title of the report job. |
| Template | Specifies the name of the template. |
| Scope | Specifies the scope of the report. |
| Scope Type | Specifies if the report is generated for a device or a fabric. |

| Field | Description |
|---|---|
| Status | Specifies the status of the report. The status messages are as follows:<br>• Success: Report is generated successfully.<br>• Scheduled: A report generating schedule is set.<br>• Running: A report job is running.<br>• Failed: Report execution failed for one or more selected switches/fabrics or an issue occurred during running of the report job.<br>• Unknown: Job state could not be identified. |
| Start At | Specifies the time at which the report is scheduled to run. |
| Last Run Time | Specifies the time at which the report was last generated. |
| User | Specifies the user who has initiated the report generation. |
| Recurrence | Specifies the frequency at which the reports are generated. |
| Internal | Specifies if the report is run generated by a user or by Nexus Dashboard Fabric Controller. The value is **false** if the report is generated by a user. |

You can perform the following actions in this tab:



**Note** You cannot perform these actions on internal report definitions.

| Action | Description |
|---|---|
| Edit | Allows you to edit a report.<br>**Note**<br>You cannot change the report name and template. |
| Re-run Report | Allows you to rerun a report. You can use the re-run option to generate a report before the scheduled execution time. |

| Action | Description |
|---|---|
| History | Allows you to view report job history.<br>The **Job History** window is displayed. You can view several entries per report job.<br>**Note**<br>The number of definitions displayed is defined by the following settings on **Settings > Server Settings > Reports** tab. Based on these values, the reports and history is purged.<br>• Max number of history across report definition<br>• Max number of reports per report definition |
| Delete | Allows you to delete a report job. |

# Reports

The **Reports** tab displays the reports which are run by a user.

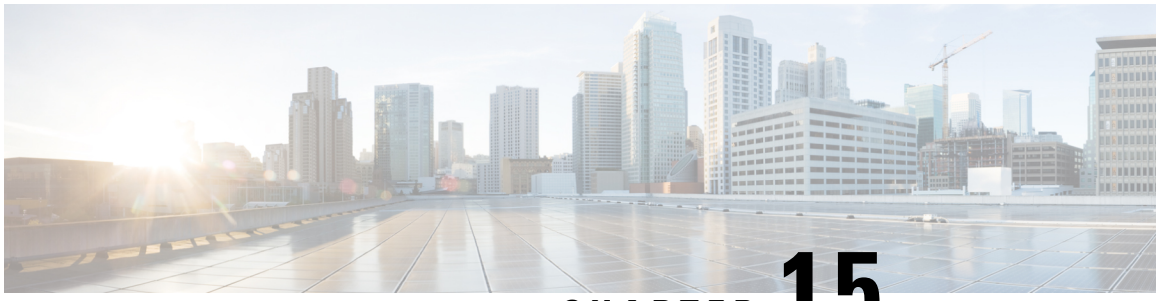
You can view the following information in this tab:

| Field | Description |
|---|---|
| Title | Specifies the title of the report.<br>• Single click on the report title opens a slide in summary panel.<br>• Double click on the report title opens the Details and Commands window. |
| Template | Specifies the name of the template. |
| Scope | Specifies the scope of the report. |
| Scope Type | Specifies if the report is generated for a device or a fabric. |
| Tag | Specifies the report is scheduled to run, or re-run report. |

| Field | Description |
|---|---|
| Status | Specifies the status of the report. The status messages are as follows:<br>• COMPLETED<br>• SUCCESS<br>• RUNNING<br>• FAILED<br>• WARNING<br>• SCHEDULED<br>• UNKNOWN |
| User | Specifies the user who has initiated the report generation. |
| Recurrence | Specifies the frequency at which the reports are generated. |
| Created At | Specifies when the report is created. |
| Internal | Specifies if the report was created by a user or Nexus Dashboard Fabric Controller. The value is false if the report is created by a user. |

You can perform the following actions in this tab:

| Action | Description |
|---|---|
| Delete | Allows you to delete a report.<br>**Note**<br>You cannot delete internal reports. |
| Compare (2 Reports) | Allows you to compare two reports side by side. The report detail is logically grouped into sections.<br>The commands are displayed based on the templates and the API that is used to run the commands on the device. For example, in the switch_inventory template, the show version, show inventory and show license usage commands are run to retrieve information. Note that the commands are displayed only if the show_and_store API is used to run the commands on the device. |
| Download | Allows you to download a report. You cannot choose more than one report to download. |



CHAPTER 15

# License Management

Beginning with Cisco Nexus Dashboard Fabric Controller Release 12.0.1a, support is removed for the following:

- Eval license state is not supported.
- Server License files are not supported.

  You must convert existing server license files to smart licenses on Cisco Smart Software Manager (CSSM). For more information, see Cisco Smart Software Manager

This chapter contains the following topics:


- Overview, on page 403
- NDFC Server Licenses, on page 404
- Smart Licensing, on page 405
- Switch Licenses, on page 408
- Switch License Files, on page 410


## Overview

You can view the existing Cisco Nexus Dashboard Fabric Controller licenses by choosing **Operations** > **License Management** > **Overview**. You can view and assign licenses in the following tabs:

- **NDFC**
- **Smart**
- **Switch License Files**

**Note** By default, the **Overview** tab appears.

The **Overview** tab has three cards namely NDFC, Switch, and Smart. These cards display the total number of licenses to purchase and the total number of licenses expiring.

To enable Smart Licensing on switches, click **Setup Smart Licensing**. For more information on Smart Licensing, check Smart section.

# NDFC Server Licenses

On NDFC tab, you can review the status of NDFC licenses for each switch. These license may be provisioned on the device, or a Smart License, or an Honor License or Unlicensed device.

Choose one or multiple switches, click **Actions** > **Assign** or **Assign All**.

When you assign a license to a device, the NDFC license service assigns the available license, based on availability on the device, status of smart licensing, and other factors.

Server based smart license is supported for Cisco MDS switches, and Nexus 9000, 3000 7000, and 5000 series of switches.

To add license from your local directory:

1. Click **Add license**.

   The **Add License File** window appears.

2. Click **Select License File** and choose appropriate files from your local directory.

3. Click **Upload** and click **Refresh** icon to refresh table and to view uploaded license files.

   The license filename, type of license, and expiration date details are extracted from the imported license file and listed in the table.

The following table displays the fields that appear on **License Management** > **NDFC**.

| Field | Description |
|---|---|
| Switch Name | Displays the name of the switch. |
| License Type | Displays the license type of the switch that can be one of the following:<br>• Switch<br>• Smart<br>• Switch Smart |
| State | Displays the license type of the switch that can be one of the following:<br>• Permanent<br>• Unlicensed<br>• Smart<br>• Expired<br>• Not Applicable<br>• Invalid |
| Expiration Date | Specifies the expiration date of license. |
| WWN/Chassis ID | Displays the world wide name or Chassis ID. |

| Field | Description |
|---|---|
| Model | Displays the model of the device. For example, DS-C9124 or N5K-C5020P-BF. |
| Fabric | Specifies the name of the fabric. |

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **License Management** > **NDFC**.

| Action Item | Description |
|---|---|
| Assign | Choose a switch, from the **Actions** drop-down list, select **Assign**.<br>A confirmation message appears. |
| Unassign | Choose a switch, from the **Actions** drop-down list, select **UnAssign**.<br>A confirmation message appears. |
| Assign All | • To assign license to all switches in the table, from the **Actions** drop-down list, choose **Assign All**.<br>A confirmation message appears<br>• Click **OK** to refresh table. |
| Unassign All | • To unassign license to all switches in the table, from the **Actions** drop-down list, choose **UnAssign All**.<br>A confirmation message appears<br>• Click **OK** to refresh table. |

# Smart Licensing

Cisco Nexus Dashboard Fabric Controller allows you to configure Smart Licensing and you can use the Smart Licensing feature to manage licenses at device-level and renew them if required.

### Introduction to Smart Licensing

Cisco Smart Licensing is a flexible licensing model that provides you with an easier, faster, and more consistent way to purchase and manage software across the Cisco portfolio and across your organization. And it's secure – you control what users can access. With Smart Licensing you get:

- **Easy Activation**: Smart Licensing establishes a pool of software licenses that can be used across the entire organization—no more PAKs (Product Activation Keys).
- **Unified Management**: My Cisco Entitlements (MCE) provides a complete view into all of your Cisco products and services in an easy-to-use portal, so you always know what you have and what you are using.
- **License Flexibility**: Your software is not node-locked to your hardware, so you can easily use and transfer licenses as needed.

To use Smart Licensing, you must first set up a Smart Account on Cisco Software Central (https://software.cisco.com/software/csws/ws/platform/home).

For a more detailed overview on Cisco Licensing, go to https://www.cisco.com/c/en/us/buy/licensing/licensing-guide.html.

### Smart License Management

This policy runs in the license microservice and provides the ability to manage the licenses for NDFC using CSSM. From this release, you can register smart licensing OnPrem, or offline mode.

When you register smart licensing on NDFC directly with internet access, Cisco Nexus Dashboard uses IP addresses to access the smart license instead of hostname and displays an error.

Ensure that subnets for IP addresses on https://smartreceiver.cisco.com are added to the routing IP address in Cisco Nexus Dashboard.

To add IP addresses, On Cisco Nexus Dashboard Web UI, navigate to **Admin Console**, **Infrastructure** > **Cluster Configuration** > **Routes** area. Click the edit icon and add IP addresses for **Management Network Routes**. Click **Save** to confirm.

The **Smart** page shows the following cards:

- **Enable Smart Licensing**

  Use the toggle switch to enable Smart Licensing. Once enabled, Smart License can assign in two ways, **Establish Trust** or **Offline Mode**.

- **Trust Status**

  Click on **Establish Trust** to establish trust. You can view two options **Transport Gateway - OnPrem with CSLU** and connecting through CSSM either directly with Cisco's licensing servers or **Proxy - Proxy via intermediate HTTP or HTTPS proxy**.

  On the **Establish Trust for Smart License** window, select the transport type to use when establishing trust with the Smart Licensing agent.

  - Choose **Default** to communicate directly with the Cisco Licensing Server.
  - Choose **Transport Gateway - OnPrem with CSLU** and enter appropriate URL.

    You don't require trust token to enable licensing. The trust is established between CSSM and the OnPrem CSLU. From NDFC and OnPrem CSLU, trust is constant, as it expected to be a local connection.
  - Choose **Proxy - Proxy via intermediate HTTP or HTTPS proxy** to transport using the proxy server. Enter the URL and Port details to access via the proxy server. For more information, refer to Smart Licensing using Policy to Establish Trust with CSSM, on page 409.

  If using the default Transport, enter the registration token obtained from CSSM.

**Note** After Smart Licensing is registered, you must manually assign licenses to the existing switches. For all switches discovered after registration, smart licenses are automatically assigned to the switches.

- **Offline Mode**

In Offline mode you can share data in alternative between NDFC instance and CSSM. Operating in an air-gap or disconnected environment, use of Offline mode allows you to export state, upload it to CSSM, and import a response back to NDFC.

To export license data and to import the response from CSSM, follow below perform the following steps:

1. On **Trust Status** click **Switch to Offline mode** to enable offline mode.
2. In offline mode with one or moe licenses assigned, click **Export License Data**.
3. On https://software.cisco.com/software/csws/ws/platform/home, navigate to smart licensing section, click **Reports** tab, and choose subsequent usage data files tab. The usage report from NDFC can be uploaded and after few minutes a response can be downloaded and imported to NDFC.
4. Click **Import License Data** and upload the CSSM acknowledge file on NDFC.

- **License Status**

  Specifies the status of the licensing on NDFC. The value is **UNCONFIGURED** if the smart licensing is not enabled. After you enable the smart licensing without registering, the value is set to **NO LICENSES IN USE**. The value is set to **IN USE** or **NOT IN USE** after registering and assigning licenses. Click the license status to view the last action, last authorization attempt, next authorization attempt, and the authorization expiry in the **License Authorization Details** pop-up window.

  Click **Policy Details**, to view smart license policy details. You can view the default smart license policy of initial 90 days and ongoing reporting within 365 days of that report.

**Note** You can view reports after 30 days of initial registering.

**Resync**

When the total number of NDFC license are not same as CSSM license counts, click **Resync**, to refresh the license counts.

Resync causes a local audit of the NDFC licenses in the switch inventory and updates the smart license counts for reporting

CSSM allows you to convert traditional licenses to Smart Licenses. For instructions, refer to https://www.cisco.com/c/dam/en/us/products/se/2020/8/Collateral/brownfield-conversion-qrg.pdf.

To migrate from Smart Licensing to Smart Licensing using Policy, launch Cisco Nexus Dashboard Fabric Controller. On the Web UI, choose **Operations > License Management > Smart** tab. Establish trust with CCSM using SLP. For instructions, refer to Smart Licensing using Policy to Establish Trust with CSSM, on page 409.

The following table describes the fields that appear in the **Switch Licenses** section.

| Field | Description |
|---|---|
| Name | Specifies the license name. |
| Count | Specifies the number of licenses used. |
| Status | Specifies the status of the licenses used. Valid values are **IN USE** and **NOT IN USE**. |

| Field | Description |
|---|---|
| Description | Specifies the type and details of the license. |

To upload or download license reports, go to https://software.cisco.com/, navigate to **Smart Software Licensing > Reports**. On **Usage Data Files** tab, click **Upload Usage Data** to upload Usage Report from NDFC. After few minutes of uploading the report, click **Download** in the **Acknowledgment** column to download a response brought back to the NDFC and imported.

# Switch Licenses

If the switch is pre-configured with a smart license, Nexus Dashboard Fabric Controller validates and assigns a switch smart license. To assign licenses to switch using the Cisco Nexus Dashboard Fabric Controller UI, choose **Operations** > **License Management** > **Smart**. Click **Enable Smart Licensing** toggle button to enable smart licensing feature.

Switch based smart license is supported for MDS switches, and Nexus 9000, and 3000 Series of switches.



**Note** For the switches which are in managed mode, switch smart license must be assigned through Nexus Dashboard Fabric Controller.

To enable switch smart license on Nexus Dashboard Fabric Controller:

- Enable smart license feature on the switch, using freeform CLI configuration.
- Configure smart licensing on the switch, using feature license smart or license smart enable command on the switch.
- Push token of your device to smart account using license smart register id token command. Use **EXEC** option in Nexus Dashboard Fabric Controller to push token.

Click **Refresh** icon to refresh table.

The following table displays the fields that appear on **License Management** > **Switch**.

| Field | Description |
|---|---|
| Switch | Displays the name of the switch. |
| Features | Displays the features on the switch. |
| Status | Displays the status of switch is in use or not.<br>• Unused<br>• In Use<br>• Out Of Compliance |

| Field | Description |
|---|---|
| Type | Displays the license type of the switch that can be one of the following:<br>• Temporary<br>• Permanent<br>• Smart<br>• Counter Permanent<br>• Unlicensed<br>• Counted |
| Warnings | Specifies the warnings about license, such as expiration date and time. |
| Group | Specifies the fabric or LAN name. |

# Smart Licensing using Policy to Establish Trust with CSSM

To establish trust with CSSM using the Smart Licensing using Policy on Cisco Nexus Dashboard Fabric Controller, perform the following steps:

### Before you begin

- Ensure that there is network reachability between Cisco Nexus Dashboard and CSSM. To configure network reachability, launch **Cisco Nexus Dashboard Web UI**. On **Admin Console**, choose **Infrastructure** > **Cluster Configuration** > **General** tab. In **Routes** area, click the edit icon, and add IP addresses for Data Network Routes. Click **Save** to confirm.
- Ensure that you have obtained the Token from CSSM.

### Procedure

**Step 1** Choose **Operations** > **License Management** > **Smart** tab.

**Step 2** Use the **Enable Smart Licensing** toggle button to enable smart licensing.

**Step 3** On the **Trust Status** card, click **Establish Trust**.

The **Establish Trust for Smart License** window appears.

**Step 4** Select the **Transport** option to register Smart License Agent.

The options are:

- **Default - NDFC communicates directly with Cisco's licensing servers**

  This option uses the following URL: https://smartreceiver.cisco.com/licservice/license.

- **Transport Gateway – OnPrem with CSLU option**

  Enter the CSLU transport URL.

**Note**
You must configure the license smart URL on the product to use the CSLU transport URL.

- **Proxy - Proxy via intermediate HTTP or HTTPS proxy**

  Enter the URL and the port if you select this option.

**Step 5** In the **Token** field, paste the token that you have obtained from CSSM to establish trust for Smart Licenses.

**Step 6** Click **Establish Trust**.

A message appears as confirmation.

The status changes from UNTRUSTED to TRUSTED. The name, count, and status of switch licenses appear.

Click on **TRUSTED** to see the details. The switch details are updated under the Switches/VDCs section of the License Assignments tab. The license type and the license state of switches that are licensed using the smart license option are Smart.

**Step 7** Click **NDFC** tab.

**Step 8** From the Actions drop-down list, select **Assign All**.

The **Status** of the server licenses shows **InCompliance**.

If the status shows **OutofComplaince**, visit the CSSM portal to acquire the required licenses.

For all other statuses, contact Cisco Technical Assistance Center (TAC).

# Switch License Files

Cisco Nexus Dashboard Fabric Controller allows you to upload multiple licenses at a single instance. Nexus Dashboard Fabric Controller parses the license files and extract the switch serial numbers. It maps the serial numbers in the license files with the discovered fabric to install the licenses on each switch. License files are moved to bootflash and installed.

The following table describes the fields that appear on this tab.

| Field | Description |
|---|---|
| Switch | Specifies the switch name. |
| Switch IP | Specifies the switch IP address. |
| License File | Specifies the type of license file. |
| Status | Specifies the status of license. |
| Result Message | Specifies the license details. |
| Last Upload Time | Specifies the date and time uploaded on server. |
| Features | Specifies the license features. |

# Adding Switch License Files

To bulk install licenses to the switches on the Cisco Nexus Dashboard Fabric Controller Web Client UI, perform the following steps:

**Procedure**

**Step 1** Choose **Operations** > **License Management** > **Switch License Files**.

The **Switch License File** window appears.

**Step 2** On the Switch License File tab, click **Add License** to upload the appropriate license file.

The **Add License File** window appears.

**Step 3** In the Add License File, click **Select License File**.

Navigate and choose the appropriate license file located in your local directory.

**Step 4** Click **Upload**.

The License file is uploaded to the Nexus Dashboard Fabric Controller. The following information is extracted from the license file.

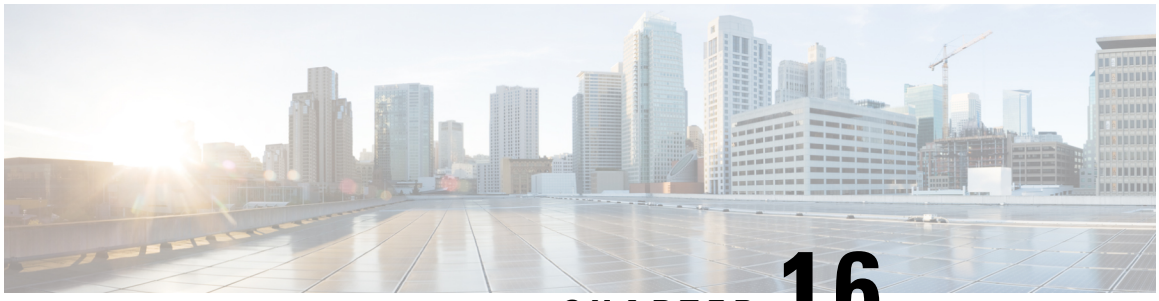
- Switch IP – IP Address of the switch to which this license is assigned.
- License File – filename of the license file
- Features List –list of features supported by the license file

**Step 5** Select the set of licenses that you want to upload and install on their respective switches. A license file is applicable for a single specific switch.

**Step 6** Click **Actions** > **Install** to install licenses.

The selected licenses are uploaded and installed on their respective switches. Status messages, including any issues or errors are updated for each file as it completes.

**Step 7** After the license matches with respective devices and installs, the **Status** column displays the status.



CHAPTER 16

# Templates


- Templates, on page 413
- Template Usage, on page 442


# Templates

**UI Navigation**

- Choose **Operations** > **Templates**.

You can add, edit, or delete templates that are configured across different Cisco Nexus, IOS-XE, IOS-XR, and Cisco MDS platforms using Cisco Nexus Dashboard Fabric Controller Web client. The following parameters are displayed for each template that is configured on Cisco Nexus Dashboard Fabric Controller 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.

***Table 46: Template Table Fields and Description***

| Field | Description |
|---|---|
| Name | Specifies the template name. |
| Supported Platforms | Specifies the platforms that the template support. |
| Type | Specifies the template type. |
| Sub Type | Specifies the template sub type. |
| Modified | Specifies the date and time of the template modification. |
| Tags | Specifies if the template is tagged to a fabric or a device. |
| Description | Specifies the template description. |
| Reference Count | Specifies the number of times a template is used. |

Click the table header to sort the entries in alphabetical order of that parameter.



**Note** Templates with errors are not listed in the Templates window. You cannot import templates with errors. To import such templates, fix the errors, and import them.

The following table describes the action items, in the **Actions** drop-down list, that appears in the **Templates** window.

***Table 47: Templates Actions and Description***

| Actions | Description |
|---|---|
| Create new template | Allows you to create a new template. For more information, see Creating a New Template, on page 415. |
| Edit template properties | Allows you to edit the template properties. You can edit only one template at a time. For more information, see Editing a Template, on page 417. |
| Edit template content | Allows you to edit the template content. You can edit only one template at a time. For more information, see Editing a Template, on page 417. |
| Duplicate template | Allows you to duplicate the selected template with a different name. You can edit the template as required. You can duplicate only one template at a time.<br><br>To duplicate a template, select the check box next to the template that you want to duplicate and choose **Duplicate template**. The **Duplicate Template** window appears. Specify a name for the duplicated template. For more information about editing the duplicated template, see Editing a Template, on page 417. |

| Actions | Description |
|---|---|
| Delete template | Allows you to delete a template. You can delete more than one template in a single instance.<br>You can delete the user-defined templates. However, you cannot delete the predefined templates<br>To delete a template, select the check box next to the template that you want to delete and choose **Delete template**. A warning message appears. If you are sure you want to delete the template, click **Confirm**. If not, click **Cancel**. If the template is in use or is a shipping template, you cannot delete it, and an error message appears.<br>**Note**<br>Select multiple templates to delete them at the same instance.<br>To delete the template permanently, delete the template that is located in your local directory: `Cisco Systems\dcn\ndfc\data\templates\`. |
| Import | Allows you to import a template from your local directory, one at a time. For more information, see Importing a Template, on page 418. |
| Import as Zip | Allows you to import `.zip` file, that contains more than one template that is bundled in a `.zip` format<br>All the templates in the ZIP file are extracted and listed in the table as individual templates.<br>For more information, see Importing a Template, on page 418. |
| Export | Allows you to export the template configuration to a local directory location. You can export only one template at a time.<br>To export a template, use the check box next to it to select it and choose **Export**. Select a location on your local system directory to store the template file. Click **Save**. The template file is exported to your local directory. |

You can only view templates with the **network-operator** role. You cannot create, edit, or save templates with this role. However, you can create or edit templates with the **network-stager** role.

This section contains the following:

# Creating a New Template

**Nexus Dashboard Fabric Controller UI Navigation**

- Choose **Operations** > **Templates**.

To create user-defined templates and schedule jobs from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Procedure

**Step 1** In the **Templates** window, from the **Actions** drop-down list, choose **Create new template**.

The **Create Template** window appears.

**Step 2** In the **Template Properties** page of the window, specify a template name, description, tags, and choose supported platforms for the new template. Next, choose a template type and a sub template type from the drop-down lists. Choose a content type for the template from the drop-down list.

**Note**
The base templates are CLI templates.

**Step 3** Click **Next** to continue editing the template or click **Cancel** to discard the changes.

The edited template properties are displayed in the **Template Content** page of the **Edit Template** window. For information about the structure of the Configuration Template, see the *Template Structure* section.

**Step 4** Click **Validate** to validate the template syntax.

**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 5** Click **Help** to open the **Editor Help** pane on the right.

This window contains more information about the format, variables, content and data types used to build the template. Close the **Editor Help** pane.

**Step 6** Click **Errors** and **Warnings** if the links are displayed. If there are no errors or warnings, the links are not available. If errors or warnings are present, and you click the links, the **Errors & Warnings** pane appears on the right displaying the errors and warnings. Close the **Errors & Warnings** pane.

**Step 7** To build the template content, select the required theme, key binding, and font size from the drop-down list.

**Step 8** Click **Finish** to complete editing of the template, click **Cancel** to discard the changes, click **Previous** to go to the **Template Properties** page.

The page with the message that the template was created appears. The page also displays the template name, type, and sub type, and the platforms. You can also click **Create another template** to create one more template or click **Edit <template name> template** to edit the template that was just edited.

**Step 9** Close the **Edit Template** window or Click **Back to template library** to go back to the **Templates** window.

# Editing a Template

**Nexus Dashboard Fabric Controller UI Navigation**

- Choose **Operations** > **Templates**.

You can edit the user-defined templates. However, the predefined templates and templates that are already published cannot be edited.

Use the **Edit Template** window to first edit the template properties and then edit the template content. Furthermore, you can edit either only the template properties using the **Edit template properties** action or only the template content using the **Edit template content** action. In other words, you can edit the template properties at one instance, and then, edit the template content at another instance. You can also use this window to view the template properties and content.

Perform the following steps to edit the template properties and then edit the template content:

## Procedure

**Step 1** In the **Templates** window, select a template. From the **Actions** drop-down list, choose **Edit template properties**.

The **Edit Template** window appears.

**Step 2** In the **Template Properties** page of the window displays the name of the template along with its description, supported platforms, tags, and content type. You can edit the template description and tags. To edit the supported platforms, clear the selected check boxes to select other switches. Next, choose a template type and a sub template type from the drop-down lists.

**Step 3** Click **Next** to continue editing the template or click **Cancel** to discard the changes.

The edited template properties are displayed in the **Template Content** page of the **Edit Template** window.

**Step 4** Click **Validate** to validate the template syntax.

**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 5** Click **Help** to open the **Editor Help** pane on the right.

This window contains more information about the format, variables, content and data types used to build the template. Close the **Editor Help** pane.

**Step 6** Click **Errors** and **Warnings** if the links are displayed. If there are no errors or warnings, the links are not available. If errors or warnings are present, and you click the links, the **Errors & Warnings** pane appears on the right displaying the errors and warnings. Close the **Errors & Warnings** pane.

**Step 7** To build the template content, select the required theme, key binding, and font size from the drop-down list.

**Step 8** Click **Finish** to complete editing of the template, click **Cancel** to discard the changes, click **Previous** to go to the **Template Properties** page.

The page with the message that the template is saved appears. The page also displays the template name, type, and sub type, and the platforms. You can also click **Create another template** to create one more template or click **Edit <template name> template** to edit the template that was just edited.

**Step 9** Close the **Edit Template** window or Click **Back to template library** to go back to the **Templates** window.

# Importing a Template

**Nexus Dashboard Fabric Controller UI Navigation**

- Choose **Operations** > **Templates**.

Follow the same procedure while importing zipped templates.



**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.

To import a template from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** In the **Templates** window, from the **Actions** drop-down list, choose **Import template**.

The **Import Template** window appears.

**Step 2** Browse and select the template that is saved on your computer.

**Step 3** Click **OK** to import the template or click **Cancel** to discard the template.

**Note**
After importing a zipped template file, either a successful or error message appears. Click **OK**.

**Step 4** You can edit the template parameters and content, if necessary. For more information, see Editing a Template, on page 417.

**Note**
When importing a zipped template file, the **Edit Template** window may not appear. However, you can edit the template parameters and content, if necessary, using the **Edit Template** action.

**Step 5** If you do not want to edit the template properties or content, then keep clicking **Next**, then **Finish** and **Back to template library** to go back to the **Templates** window.

# 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, N6K, N7K, N9K, MDS, VDC, N9K-9000v, IOS-XE, IOS-XR, Others, All Nexus Switches list separated by comma. | No |
| templateType | Specifies the type of Template used. | • CLI<br>• POAP<br>**Note** POAP option is not applicable for Cisco Nexus Dashboard Fabric Controller LAN Fabric deployment.<br>• POLICY<br>• SHOW<br>• PROFILE<br>• FABRIC<br>• ABSTRACT<br>• REPORT | Yes |

| Property Name | Description | Valid Values | Optional? |
|---|---|---|---|
| templateSubType | Specifies the sub type associated with the template. | | |

| Property Name | Description | Valid Values | Optional? |
|---|---|---|---|
| | | • CLI<br>  • N/A<br>• POAP<br>  • N/A<br>  • VXLAN<br>  • FABRICPATH<br>  • VLAN<br>  • PMN<br>**Note** POAP option is not applicable for Cisco Nexus Dashboard Fabric Controller LAN Fabric deployment.<br>• POLICY<br>  • VLAN<br>  • INTERFACE_VLAN<br>  • INTERFACE_VPC<br>  • INTERFACE_ETHERNET<br>  • INTERFACE_BD<br>  • INTERFACE_PORT_CHANNEL<br>  • INTERFACE_FC<br>  • INTERFACE_MGMT<br>  • INTERFACE_LOOPBACK<br>  • INTERFACE_NVE<br>  • INTERFACE_VFC<br>  • INTERFACE_SUBINTERFACE_PORT_CHANNEL<br>  • DEVICE<br>  • FEX<br>  • INTRA_FABRIC_LINK<br>  • INTER_FABRIC_LINK | |

| Property Name | Description | Valid Values | Optional? |
|---|---|---|---|
| | | • INTERFACE<br>• SHOW<br>• VLAN<br>• INTERFACE_VLAN<br>• INTERFACE_VPC<br>• INTERFACE_ETHERNET<br>• INTERFACE_BD<br>• INTERFACE_PORTCHANNEL<br>• INTERFACE_FC<br>• INTERFACE_MGMT<br>• INTERFACE_LOOPBACK<br>• INTERFACE_NVE<br>• INTERFACE_VFC<br>• INTERFACE_SAN_PORTCHANNEL<br>• DEVICE<br>• FEX<br>• INTRA_FABRIC_LINK<br>• INTER_FABRIC_LINK<br>• INTERFACE<br>• PROFILE<br>• VXLAN<br>• FABRIC<br>• NA | |

| Property Name | Description | Valid Values | Optional? |
|---|---|---|---|
| | | • ABSTRACT<br>• VLAN<br>• INTERFACE_VLAN<br>• INTERFACE_VPC<br>• INTERFACE_ETHERNET<br>• INTERFACE_BD<br>• INTERFACE_PORT_CHANNEL<br>• INTERFACE_FC<br>• INTERFACE_MGMT<br>• INTERFACE_LOOPBACK<br>• INTERFACE_NVE<br>• INTERFACE_VFC<br>• INTERFACE_SAN_PORT_CHANNEL<br>• DEVICE<br>• FEX<br>• INTRA_FABRIC_LINK<br>• INTER_FABRIC_LINK<br>• INTERFACE<br>• REPORT<br>• UPGRADE<br>• GENERIC | |

| Property Name | Description | Valid Values | Optional? |
|---|---|---|---|
| contentType | | • CLI<br>• TEMPLATE_CLI<br>• POAP<br>• TEMPLATE_CLI<br>**Note** POAP option is not applicable for Cisco Nexus Dashboard Fabric Controller LAN Fabric deployment.<br>• POLICY<br>• TEMPLATE_CLI<br>• PYTHON<br>• SHOW<br>• TEMPLATE_CLI<br>• PROFILE<br>• TEMPLATE_CLI<br>• PYTHON<br>• FABRIC<br>• PYTHON<br>• ABSTRACT<br>• TEMPLATE_CLI<br>• PYTHON<br>• REPORT<br>• PYTHON | Yes |
| implements | Used to implement the abstract template. | Text | Yes |
| dependencies | Used to select the specific feature of a switch. | Text | Yes |

| Property Name | Description | Valid Values | Optional? |
|---|---|---|---|
| published | Used to Mark the template as read only and avoids changes to it. | “true” or “false” | Yes |

## Template Variables

This section contains declared variables, the data type, default values, and valid values conditions for the parameters that are used in the template. These declared variables are used for value substitution in the template content section during the dynamic command generation process. Also these variables are used in decision making and in iteration blocks in the template content section. Variables have predefined data types. You can also add a description about the variable. The following table describes the syntax and usage for the available datatypes.

| Variable Type | Valid Value | Iterative? |
|---|---|---|
| boolean | true\|false | No |
| enum | `Example: running-config, startup-config` | No |
| float | Floating number format | No |
| floatRange | `Example: 10.1,50.01` | Yes |
| Integer | Any number | No |
| integerRange | Contiguous numbers separated by “-“<br>Discrete numbers separated by “,”<br>`Example: 1-10,15,18,20` | Yes |
| interface | Format: <if type><slot>[/<sub slot>]/<port><br>`Example: eth1/1, fa10/1/2 etc.` | No |
| interfaceRange | `Example: eth10/1/20-25, eth11/1-5` | Yes |
| ipAddress | IPv4 OR IPv6 address | No |

| Variable Type | Valid Value | Iterative? |
|---|---|---|
| ipAddressList | You can have a list of IPv4, IPv6, or a combination of both types of addresses.<br><br>Example 1: 172.22.31.97,<br>172.22.31.99,<br>172.22.31.105,<br>172.22.31.109<br>Example 2:<br>2001:0db8:85a3:0000:0000:8a2e:0370:7334,<br><br>2001:0db8:85a3:0000:0000:8a2e:0370:7335,<br><br>2001:0db8:85a3:1230:0000:8a2f:0370:7334<br>Example 3: 172.22.31.97,<br>172.22.31.99,<br><br>2001:0db8:85a3:0000:0000:8a2e:0370:7334,<br><br>172.22.31.254 | Yes |
| ipAddressWithoutPrefix | Example: 192.168.1.1<br><br>or<br><br>Example: 1:2:3:4:5:6:7:8 | No |
| ipV4Address | IPv4 address | No |
| ipV4AddressWithSubnet | Example: 192.168.1.1/24 | No |
| ipV6Address | IPv6 address | No |
| ipV6AddressWithPrefix | Example: 1:2:3:4:5:6:7:8<br><br>22 | No |
| ipV6AddressWithSubnet | IPv6 Address with Subnet | No |
| ISISNetAddress | Example:<br>49.0001.00a0.c96b.c490.00 | No |
| long | Example: 100 | No |
| macAddress | 14 or 17 character length MAC address format | No |
| string | Free text, for example, used for the description of a variable<br><br>Example:<br>string scheduledTime<br>{<br><br>regularExpr=^([01]\d\|2[0-3]):([0-5]\d)$;<br>} | No |

| Variable Type | Valid Value | Iterative? |
|---|---|---|
| string[] | Example: {a,b,c,str1,str2} | Yes |
| struct | Set of parameters that are bundled under a single variable.<br><br>struct <structure name declaration > {<br><parameter type> <parameter 1>;<br><parameter type> <parameter 2>;<br>…..<br>} [<structure_inst1>] [, <structure_inst2>] [, <structure_array_inst3 []>];<br><br>struct interface_detail {<br>string inf_name;<br>string inf_description;<br>ipAddress inf_host;<br>enum duplex {<br>validValues = auto, full, half;<br>};<br>}myInterface, myInterfaceArray[]; | No<br>**Note**<br>If the struct variable is declared as an array, the variable is iterative. |
| wwn<br>(Available only in Cisco Nexus Dashboard Fabric Controller Web Client) | Example: 20:01:00:08:02:11:05:03 | No |

## Variable Meta Property

Each variable that is defined in the template variable section has a set of meta properties. The meta properties are mainly the validation rules that are defined for the variable.

The following table describes the various meta properties applicable for the available variable types.

| Variable Type | Description | Variable Meta Property | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | default Value | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| boolean | A boolean value.<br>Example: true | Yes | | | | | | | | | | | |
| enum | | | Yes | | | | | | | | | | |

| Variable Type | Description | Variable Meta Property | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | default Value | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| float | signed real number<br>Example:<br>75.56, -8.5 | Yes | Yes | Yes | Yes | Yes | | | | | | | |
| floatRange | range of signed real numbers<br>Example:<br>50.5 - 54.75 | Yes | Yes | Yes | Yes | Yes | | | | | | | |
| integer | signed number<br>Example:<br>50, -75 | Yes | Yes | | Yes | Yes | | | | | | | |
| integerRange | Range of signed numbers<br>Example:<br>50-65 | Yes | Yes | | Yes | Yes | | | | | | | |
| interface | specifies interface/port<br>Example:<br>Ethernet 5/10 | Yes | Yes | | | | Yes | Yes | Yes | Yes | | | |
| interfaceRange | | Yes | Yes | | | | Yes | Yes | Yes | Yes | | | |
| ipAddress | IP address in IPv4 or IPv6 format | Yes | | | | | | | | | | | |

| Variable Type | Description | Variable Meta Property | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | default Value | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| ipAddrList | You can have a list of IPv4, IPv6, or a combination of both types of addresses.<br>Example 1: 172.22.31.97, 172.22.31.99, 172.22.31.105, 172.22.31.109<br>Example 2: [redacted]<br>Example 3: 172.22.31.97, 172.22.31.99, [redacted], 172.22.31.254<br>**Note** Separate the addresses in the list using commas and not hyphens. | Yes | | | | | | | | | | | |

| Variable Type | Description | Variable Meta Property | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | default Value | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| ipAddrMask | IPv4 or IPv6 Address (does not require prefix/mask) | | | | | | | | | | | | |
| ipV4Addr | IPv4 address | Yes | | | | | | | | | | | |
| ipV4AddrMask | IPv4 Address with Subnet | Yes | | | | | | | | | | | |
| ipV6Addr | IPv6 address | Yes | | | | | | | | | | | |
| ipV6AddrMask | IPv6 Address with prefix | Yes | | | | | | | | | | | |
| ipV6AddrMask | IPv6 Address with Subnet | Yes | | | | | | | | | | | |
| ISNAddr | Example: 49.0000.0000.0000.00 | | | | | | | | | | | | |
| long | Example: 100 | Yes | | | Yes | Yes | | | | | | | |
| macAddr | 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>Example for string<br>Regular expression: string<br>shutdown {<br>no<br>} | Yes | | | | | | | | | Yes | Yes | Yes |
| string[] | string literals that are separated by a comma (,)<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 parameters that are bundled under a single variable.<br>`struct <structure name declaration > { <parameter type> <parameter 1>; <parameter type> <parameter 2>; ….. } [structvar1] [, structvar2] [, structvar3 []>];` | | | | | | | | | | | | |
| wwn | WWN address | | | | | | | | | | | | |

### Example: Meta Property Usage

```
##template variables

integer VLAN_ID {
min = 100;
max= 200;
};

string USER_NAME {
defaultValue = admin123;
minLength = 5;
};

struct interface_a{
```

```
 string inf_name;
 string inf_description;
 ipAddress inf_host;
 enum duplex {
  validValues = auto, full, half;
 };
}myInterface;

##
```

## Variable Annotation

You can configure the variable properties marking the variables using annotations.

**Note** Variable Annotations are available for POAP only. However, the annotations do not impact on the template type 'CLI'.

The following annotations can be used in the template variable section.

| Annotation Key | Valid Values | Description |
|---|---|---|
| AutoPopulate | Text | Copies values from one field to another |
| DataDepend | Text | |
| Description | Text | Description of the field appearing in the window |
| DisplayName | Text<br>**Note**<br>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>**Note** Use this annotation only for the ipAddress variable. |
| IsManagementIP | “true” or “false”<br>**Note** This annotation must be marked only for variable “ipAddress”. | |
| IsMandatory | “true” or “false” | Validates if a value should be passed to the field mandatorily |
| IsMTU | “true” or “false” | |
| IsMultiCastGroupAddress | “true” or “false” | |
| IsMultiLineString | “true” or “false” | Converts a string field to multiline string text area |
| IsMultiplicity | “true” or “false” | |
| IsPassword | “true” or “false” | |
| IsPositive | “true” or “false” | Checks if the value is positive |
| IsReplicationMode | “true” or “false” | |
| IsShow | “true” or “false” | Displays or hides a field on the window |
| IsSiteId | “true” or “false” | |
| IsSourceDevice | “true” or “false” | |
| IsSourceFabric | “true” or “false” | |
| IsSourceInterface | “true” or “false” | |
| IsSourceSwitchName | “true” or “false” | |

| Annotation Key | Valid Values | Description |
|---|---|---|
| IsSwitchName | “true” or “false” | |
| IsRMID | “true” or “false” | |
| IsVPCDomainID | “true” or “false” | |
| IsVPCID | “true” or “false” | |
| IsVPCPeerLinkPort | “true” or “false” | |
| IsVPCPeerLinkPortChannel | “true” or “false” | |
| IsVPCPortChannel | “true” or “false” | |
| Password | Text | Validates the password field |
| PeerOneFEXID | “true” or “false” | |
| PeerTwoFEXID | “true” or “false” | |
| PeerOnePCID | “true” or “false” | |
| PeerTwoPCID | “true” or “false” | |
| PrimaryAssociation | | |
| ReadOnly | “true” or “false” | Makes the field read-only |
| ReadOnlyOnEdit | “true” or “false” | |
| SecondaryAssociation | Text | |
| Section | | |
| UsePool | “true” or “false” | |
| UseDNSReverseLookup | | |
| Username | Text | Displays the username field on the window |
| Warning | Text | Provides text to override the Description annotation |

### Example: AutoPopulate Annotation

```
##template variables
string BGP_AS;
 @(AutoPopulate="BGP_AS")
  string SITE_ID;
##
```

### Example: DisplayName Annotation

```
##template variables
@(DisplayName="Host Name", Description = "Description of the host")
String hostname;
@(DisplayName="Host Address", Description = " test description" IsManagementIP=true)
ipAddress hostAddress;
##
```

### Example: IsMandatory Annotation

```
##template variables
@(IsMandatory="ipv6!=null")
ipV4Address ipv4;
@(IsMandatory="ipv4!=null")
ipV6Address ipv6;
##
```

### Example: IsMultiLineString Annotation

```
##template variables
@(IsMultiLineString=true)
string EXTRA_CONF_SPINE;
##
```

### IsShow Annotation

```
##template variables
boolean isVlan;
@(IsShow="isVlan==true")
integer vlanNo;
##

##template variables
boolean enableScheduledBackup;
@(IsShow="enableScheduledBackup==true",Description="Server time")
string scheduledTime;
##
The condition "enableScheduledBackup==true" evaluates to true/false

##template variables
@(Enum="Manual,Back2BackOnly,ToExternalOnly,Both")
string VRF_LITE_AUTOCONFIG;
@(IsShow="VRF_LITE_AUTOCONFIG!=Manual",  Description="Target Mask")
integer DCI_SUBNET_TARGET_MASK
##
The condition "VRF_LITE_AUTOCONFIG!=Manual" matches string comparison to evaluate to true
or false
```

### Example: Warning Annotation

```
##template variables
@(Warning="This is a warning msg")
  string SITE_ID;
##
```

# Templates Content

This section includes the configuration commands and any parameters that you want to include in the template. These commands can include the variables declared in the template variables section. During the command generation process the variable values are substituted appropriately in the template content.



**Note** You must specify the commands that you include as if you were entering them in the global configuration command mode on any device. You must consider the command mode when you include commands.

Template content is governed by the usage of variables.

- Scalar variables: does not take a range or array of values which cannot be used for iteration (In the variable types table those marked iterate-able as 'No'). Scalar variables must be defined inside the template content.

```
Syntax: $$<variable name>$$
Example: $$USER_NAME$$
```

- Iterative variables: used for block iteration. These loop variable must be accessed as shown below inside the iteration block.

```
Syntax:@<loop variable>
Example:
foreach val in $$INTEGER_RANGE_VALUE$$ {
@val
}
```

- Scalar Structure Variable: Structure member variables can be accessed inside the template content.

```
Syntax: $$<structure instance name>.<member variable name>$$
Example: $$myInterface.inf_name$$
```

- Array Structure Variable: Structure member variables can be accessed inside the template content.

```
Syntax: $$<structure instance name>.<member variable name>$$
Example: $$myInterface.inf_name$$
```

In addition to the template variables, you can use the conditional and iterative command generation using the following statements:

- if-else if-else Statement: makes a logical decision in inclusion/exclusion of set of configuration command based on the value assigned for the variable in it.

```
Syntax: if(<operand 1>  <logical operator>  <operand 2>){
command1 ..
command2..
..
}
else  if (<operand 3> <logical operator> <operand 4> )
{
Command3 ..
Command4..
..
}
else
{
```

```
Command5 ..
Command6..
..
}
Example: if-else if-else statement
if($$USER_NAME$$ == 'admin'){
Interface2/10
no shut
}
else {
Interface2/10
shut
}
```

- foreach Statement: used for iterating a block of commands. The iteration is performed based on the assigned loop variable value.

```
Syntax:
foreach <loop index variable> in $$<loop variable>$$ {
@<loop index variable> ..
}
Example: foreach Statement
foreach ports in $$MY_INF_RANGE$${
interface @ports
no shut
}
```

- Optional parameters: By default all parameters are mandatory. To make a parameter optional, you must annotate the parameter.

In the variable section, you can include the following command:

- **@(IsMandatory=false)**
- **Integer frequency;**

  In the template content section, a command can be excluded or included without using “if” condition check, by assigning a value to the parameter. The optional command can be framed as below:

- **probe icmp [frequency frequency-value] [timeout seconds] [retry-count retry-count-value]**

## Advanced Features

The following are the advanced features available to configure templates.

- Assignment Operation

  Config template supports assignment of variable values inside the template content section. The values are validated for the declared data type of the variable. If there is a mismatch, the value is not assigned.

  Assignment operation can be used under the following guidelines:

  - The operator on the left must be any of the template parameters or a for loop parameter.
  - The operator on the right values can be any of the values from template parameters, for loop parameters, literal string values surrounded by quotes or simple string values.

  If a statement does not follow these guidelines, or if it does not suit this format, it will not be considered as assignment operation. It is substituted during command generation like other normal lines.

```
Example: Template with assignment operation
##template properties
name =vlan creation;
userDefined= true;
supportedPlatforms = All;
templateType = CLI;
published = false;
##
##template variables
integerRange vlan_range;
@(internal=true)
integer vlanName;
##
##template content
foreach vlanID in $$vlan_range$${
vlan @vlanID
$$vlanName$$=@vlanID
name myvlan$$vlanName$$
}
##
```

- Evaluate methods

  Config template uses the Java runtime provided Java script environment to perform arithmetic operations (such as ADD, SUBTRACT, and so on), string manipulations, and so on.

  Locate the JavaScript file in the template repository path. This file contains primary set of arithmetic, string functions. You can also add custom JavaScript methods.

  These methods can be called from config template content section in below format:

```
Example1:
$$somevar$$ = evalscript(add, "100", $$anothervar$$)
```

  Also the *evalscript* can be called inside if conditions as below:

```
if($$range$$ > evalscript(sum, $$vlan_id$$,  -10)){
do something...
}
```

  You can call a method that is located at the backend of the Java script file.

- Dynamic decision

  Config template provides a special internal variable “LAST_CMD_RESPONSE”. This variable stores the last command response from the device during the execution of the command. This can be used in the config template content to make dynamic decisions to deliver the commands that are based on the device condition.



**Note** The if block must be followed by an else block in a new line, which can be empty.

An example use case to create a VLAN, if it is does not exist on the device.

```
Example: Create VLAN
##template content
show vlan id $$vlan_id$$
if($$LAST_CMD_RESPONSE$$ contains "not found"){
vlan $$vlan_id$$
}
```

```
else{
}
##
```

This special implicit variable can be used only in the “IF” blocks.

- Template referencing

  You can have a base template with all the variables defined. This base template can be imported to multiple templates. The base template content is substituted in the appropriate place of the extending template. The imported template parameters and the contents can be accessed inside the extending template.

```
Example: Template Referencing
Base template:
##template properties
 name =a vlan base;
 userDefined= true;
 supportedPlatforms = All;
 templateType = CLI;
 published = false;
 timestamp = 2015-07-14 16:07:52;
 imports = ;
##
##template variables
 integer vlan_id;
##
##template content
 vlan $$vlan_id$$
##

Derived Template:
##template properties
 name =a vlan extended;
 userDefined= true;
 supportedPlatforms = All;
 templateType = CLI;
 published = false;
 timestamp = 2015-07-14 16:07:52;
 imports = a vlan base,template2;
##
##template variables
 interface vlanInterface;
##
##template content
 <substitute a vlan base>
 interface $$vlanInterface$$
 <substitute a vlan base>
##
```

When you launch the extended template, the parameter inputs for the base template are also obtained. In addition, the substituted content is used for complete CLI command generation.

## Report Template

The template type of REPORT template is python, and it has two subtypes, UPGRADE and GENERIC.

### 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 Nexus Dashboard Fabric Controller 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**: Specifies the name of the pool.

**POOL RANGE**: Specifies the IP address range of the pool.

**SUBNET MASK**: Specifies the subnet mask.

**MAX ENTRIES**: Specifies the maximum number of entries that can be allocated from the pool.

**USAGE INSIDE RANGE**: Specifies the current number of entries allocated inside the pool range.

**USAGE OUTSIDE RANGE**: Specifies the current number of entries set outside the pool range.

**USAGE PERCENTAGE**: This is calculated by using the formula: (Usage Inside Range/Max Entries) *100.

Click **View Details** to display a view of resources allocated or set in each resource pool. For example, the detailed section for a SUBNET has information about the resources that have been allocated within the subnet.

**Switch Inventory Report**

This report provides a summary about the switch inventory.

Click **View Details** to display more information about the modules and licenses.

**SFP Report**

This report provides information about utilization of SFPs at a fabric and device level.



**Note** The switch inventory and SFP reports are supported only on Cisco Nexus devices.

**Troubleshooting Reports**

These reports are generated to help in troubleshooting scenarios. Currently, the **NVE VNI Counters** report is the only pre-defined troubleshooting report. Generating **NVE VNI Counters** reports involves performing periodic checks to identify the VNIs that are among the top hits based on network traffic. In a large-scale setup, we recommend limiting the report generation frequency to a minimum of 60 minutes.

**NVE VNI Counters Report**

This report collects the **show nve vni counters** command output for each VNI in the fabric.

After comparing the oldest report and the newest report, the **Summary** section shows the top-10 hit VNIs. The top hit VNIs are displayed in these categories:

- L2 or L3 VNIs for unicast traffic

- L2 or L3 VNIs for multicast traffic
- L2 only VNIs for unicast traffic
- L2 only VNIs for multicast traffic
- L3 only VNIs for unicast traffic
- L3 only VNIs for multicast traffic

The oldest report refers to the first report that is saved in the current reporting task. If you want to select a specific report as the first report against which the current report has to be compared, delete all reports that are older than the one selected so that the selected report becomes the first and oldest report.

For example, three reports were run yesterday at 8:00 a.m, 4:00 p.m. and 11:00 p.m. If you want to use the report at 11:00 p.m. as the first and oldest report for today's reporting, delete the two reports that were run yesterday at 8:00 a.m. and 4:00 p.m.

For a periodic report, the oldest report is the first report that is run at the start time of a period. For daily and weekly reports, the current report is compared against the previously generated report.

The **Summary** section displays a column-wise report with information about the total transmitted bytes and the VNIs. Use the horizontal scroll bar at the bottom of the window to display more columns.



**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.

For more information on how the reports are displayed, refer *Programmable Reports* chapter.

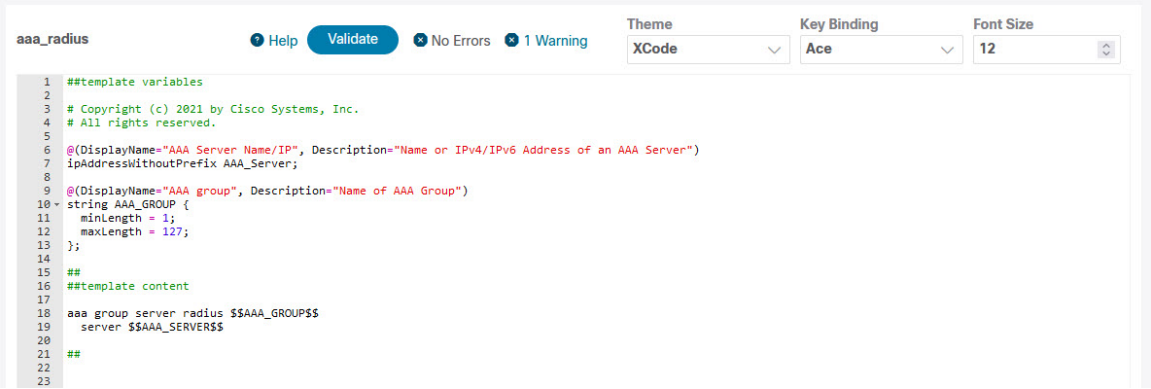
# Template Usage

| | | |
|---|---|---|
| template Type | Specifies the type of Template used. | • POLICY<br>• SHOW<br>• PROFILE<br>• FABRIC<br>• REPORT<br>• INTERNAL<br>• EXEC |

| | | |
|---|---|---|
| template content type | Specifies the type of content in the template. | • CLI<br>• PYTHON<br>• PYTHON3<br>• PYTHON3_CLI<br>• PYTHON_CLI<br>• TEXT |

# Policy Template

For the policy template, there are two template content types: CLI and PYTHON. With CLI content type, the policy templates are parameterized CLI templates. They can have a lot of variables and CLIs. Typically, CLI policy templates are small and do not have any if-else-for etc. like constructs. An example CLI policy template for AAA server configuration is shown below:




But you can also have policy templates of template content type PYTHON. Essentially, this allows multiple CLI policy templates to be combined together with a common “source” so that they get all applied/un-applied at one go. For example, when you want to create a vPC host port, it has to be created symmetrically on both peers that are part of the vPC pair. In addition, you have to create port-channel, member interfaces, channel-group, etc. This is why a python vPC host policy template has been added. An example interface PYTHON template for setting up a routed interface is shown below:

ext_int_routed_host_11_1 | Help | Validate | No Errors | No Warnings | Theme: XCode | Key Binding: Ace | Font Size: 12

```
##template variables

#    Copyright (c) 2019-2022 by Cisco Systems, Inc.
#    All rights reserved.
@(IsInternal=true)
string SERIAL_NUMBER;

@(PrimaryAssociation=true, IsInternal=true)
interface INTF_NAME;

@(IsMandatory=false, DisplayName="Interface IP", Description="IP address of the interface", ReadOnly=true)
ipV4Address IP;

@(IsMandatory="IP!=null", DisplayName="IP Netmask Length", Description="IP netmask length used with the IP address (Min:1, Max:31)", ReadOnly=true)
integer PREFIX {
  min = 1;
  max = 31;
};

@(IsMandatory=false, DisplayName="Interface IPv6", Description="IPv6 address of the interface", ReadOnly=true)
ipV6Address IPv6;

@(IsMandatory="IPv6!=null", DisplayName="IPv6 Netmask Length", Description="IPv6 netmask length used with the IPv6 address (Min:1, Max:128)", ReadOnly=true)
integer PREFIXv6 {
  min = 1;
  max = 128;
};

@(IsMandatory=false, DisplayName="Interface VRF", Description="Interface VRF name, default VRF if not specified", ReadOnly=true)
string INTF_VRF {
  minLength = 1;
  maxLength = 32;
};

@(IsMandatory=false, DisplayName="Routing TAG", Description="Routing tag associated with interface IP", ReadOnly=true)
string ROUTING_TAG;

@(DisplayName="MTU", IsMTU=true, Description="MTU for the interface", ReadOnly=true)
integer MTU {
min = 576;
max = 9216;
defaultValue=9216;
};

@(DisplayName="SPEED", Description="Interface Speed", ReadOnly=true)
enum SPEED {
validValues=Auto,100Mb,1Gb,2.5Gb,5Gb,10Gb,25Gb,40Gb,50Gb,100Gb,200Gb,400Gb;
defaultValue=Auto;
};

@(IsMandatory=false, DisplayName="Interface Description", Description="Add description to the interface", ReadOnly=true)
string DESC {
  minLength = 1;
  maxLength = 254;
};

@(IsMandatory=false, IsMultiLineString=true, DisplayName="Freeform Config", Description="Additional CLI for the interface", ReadOnly=true)
string CONF;

@(DisplayName="Enable Interface", Description="Uncheck to disable the interface", ReadOnly=true)
boolean ADMIN_STATE {
defaultValue=true;
};

@(IsInternal=true)
string SOURCE;

##
##template content

from com.cisco.dcbu.vinci.rest.services.jython import PTIWrapper
from com.cisco.dcbu.vinci.rest.services.jython import Wrapper
from com.cisco.dcbu.vinci.rest.services.jython import WrappersResp
from utility import *

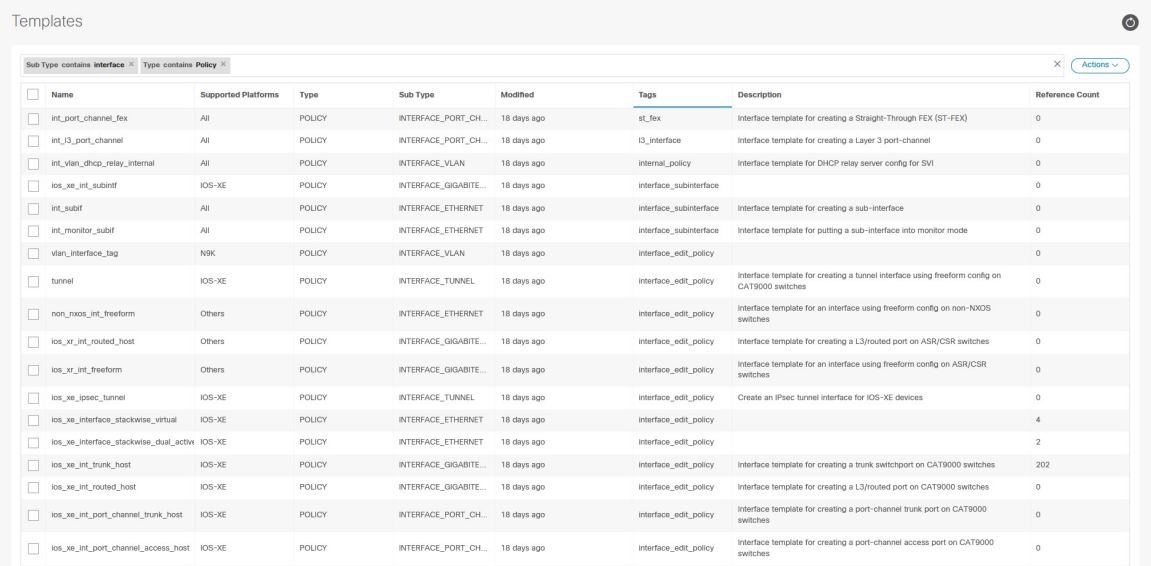
def add():
    try:

        respObj = WrappersResp.getRespObj()
        try:
            adminState = ADMIN_STATE
        except:
            adminState = "true"
            pass
        try:
            source = SOURCE
        except:
            source = INTF_NAME
            pass
        Wrapper.print("ext_int_routed_host_11_1_add : Source sn = %s, "
                      "source interface= %s: source: %s"
                      % (SERIAL_NUMBER, INTF_NAME, source))

        routingTag = ""
        try:
            if ROUTING_TAG != "":
                routingTag = ROUTING_TAG
        except:
            pass

        #Only valid operation is shut/no-shut from interface page
        #In addition, this can only happen if someone does a save on the interface edit for an interface attached to this policy
        #After that shut/no-shut from interface manager starts sending source = INTF instead of source = LINK-UUID of VRF_LITE IFC
        #This is a bug that needs to be fixed but right now putting a workaround here
        if source == INTF_NAME:
```

Each policy template has a template subtype like DEVICE, INTERFACE, etc. This allows the right policy template to appear at the right selection point. For example, in the Interface window, you will only see the interface policy templates.



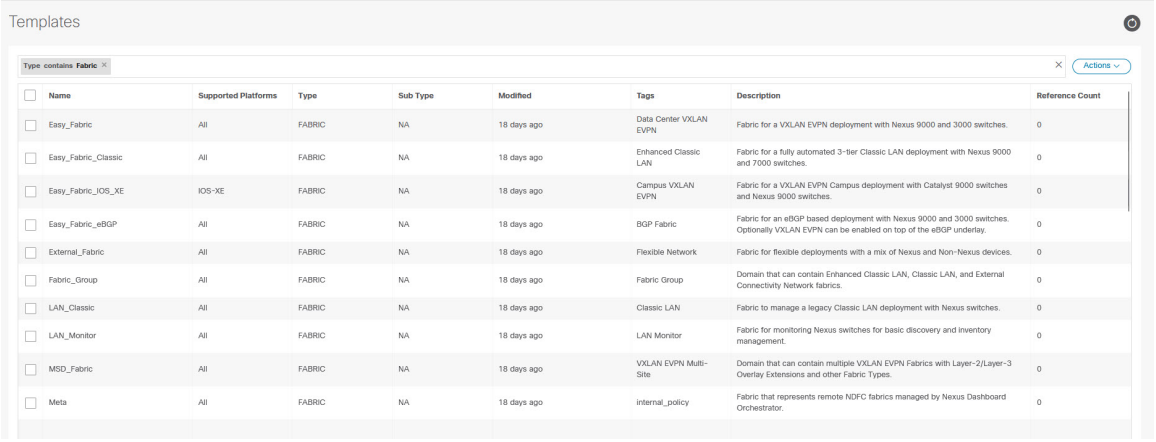

You can make a copy of any of these templates and customize them as per their needs. That is the typical use-case for customization. **Do not** modify existing policies but make a copy, and then customize as per the requirements. Otherwise, after a DCNM upgrade, the changes may be lost.

In general, a template already in use, meaning one that is already applied to some switch within any fabric, cannot be edited.

**Note** No Type-CLI templates are used in the Fabric Controller persona only. They are all replaced with more powerful Policy templates which are a super set.

# Fabric Template

A fabric template is basically a python template, specifically jython, which is java + python. A fabric template is quite comprehensive, and in that it embeds the rules that are required for deploying a fabric, including all the logic required to generate intended configuration of all switches within the entire fabric. Configuration is generated based on published Cisco best practice guidelines. In addition to the embedded rules, the fabric template also integrates with other entities such as resource manager, topology database, device roles, configuration compliance, etc. and generates the configuration accordingly for all the devices in the fabric. This is the inherent part of NDFC fabrics.



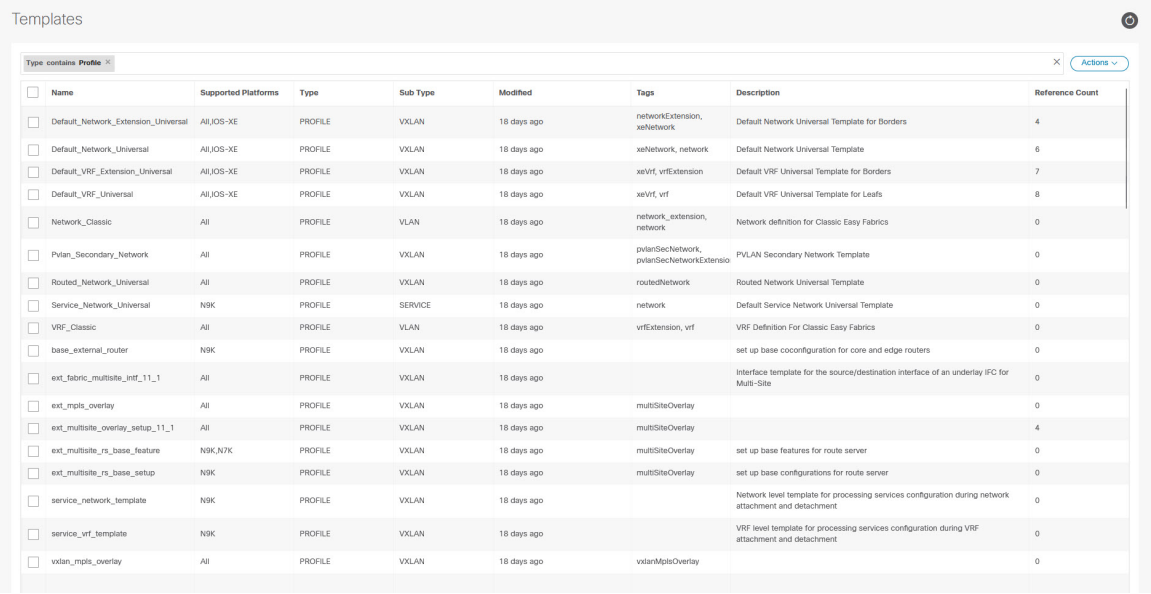

The expectation is that users will not create their own fabric templates. NDFC provides a few fabric templates out of the box such as Easy Fabric, External Fabric, MSD Fabric, eBGP Fabric, and so on.

# Profile Template

A profile template is used for provisioning of overlays (networks or VRFs). The idea is that when you apply some overlay configuration, there are multiple pieces of configurations that should go together. For example, valid layer-3 network configuration in a VXLAN EVPN fabric requires VLAN, SVI, int nve config, EVPN route-target, etc. All of these pieces are put together into what is called a configuration profile (NX-OS construct) and then effectively applied at one go. Either the whole configuration profile gets applied or nothing gets applied, on the switch. In this way, you are not left with any dangling or stray configurations on the switches. For any kind of overlay configurations, whether it is on the leaf or on the borders, NDFC employs profile templates.

There are four kinds of profile templates that are distinguished with tags as depicted below:

- Network Profile (applied to all devices with role leaf)
- Network Extension Profile (applied to all devices with role ‘border*’)
- VRF Profile (applied to all devices with role leaf)
- VRF Extension Profile (applied to all devices with role ‘border*’)




For more information about how to apply overlay configuration via the Networks & VRFs workflow in NDFC, see *Creating and Deploying Networks and VRFs* section.

### Additional Notes

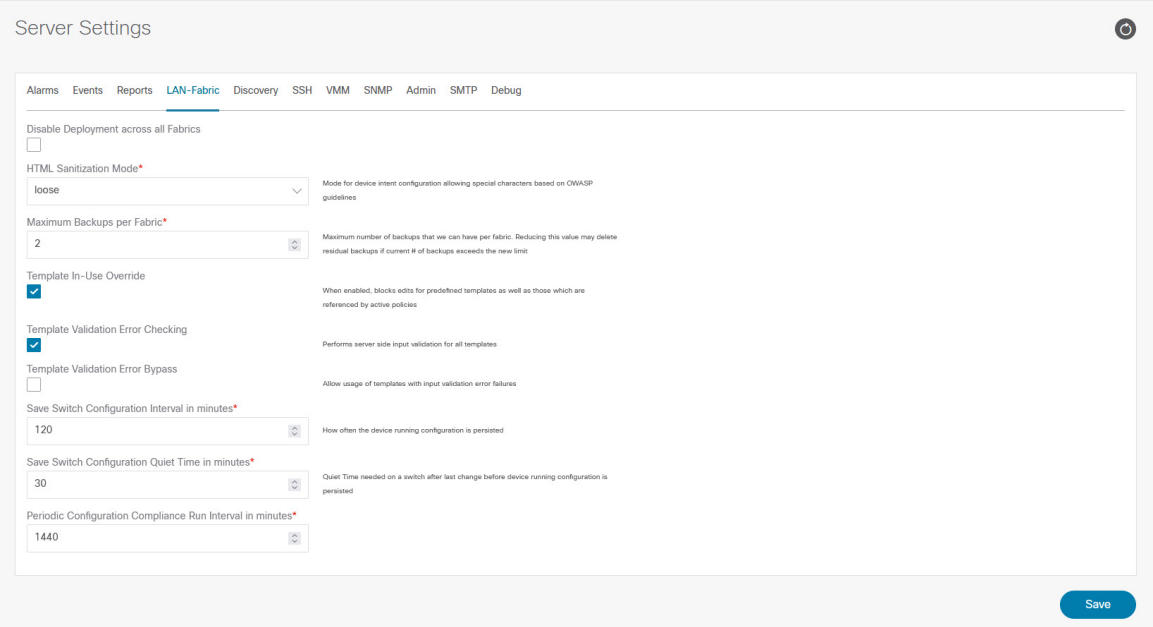
When a policy or profile template is applied, an instance is created for each application of the template. The common terminology used for this is Policy Template Instance or PTI. A PTI is effectively a policy or profile template + the Name-value pairs that give it a specific instance, post substitution. PTIs created for a device can be viewed under the View/Edit policies option for that device in Fabric Builder. In the tabular view, the View/Edit policies button allows selection and bulk creation/deletion of policies across a subset of devices in the entire fabric. For more information, see *Viewing and Editing Policies* section.

# Changing the Contents of a Template in Use

A template in general, whether it is a policy, fabric or profile template, cannot be modified once it has been instantiated. However, there could be cases where you want to edit the content of a template, like fixing a bug in the template or changing an already deployed config. This can be achieved by toggling the **Template In-Use Override** option in **Settings > Server Settings > LAN-Fabric** tab.

### Procedure

**Step 1** Under the **LAN-Fabric** tab in **Server Settings**, check the **Template In-Use Override** check box.



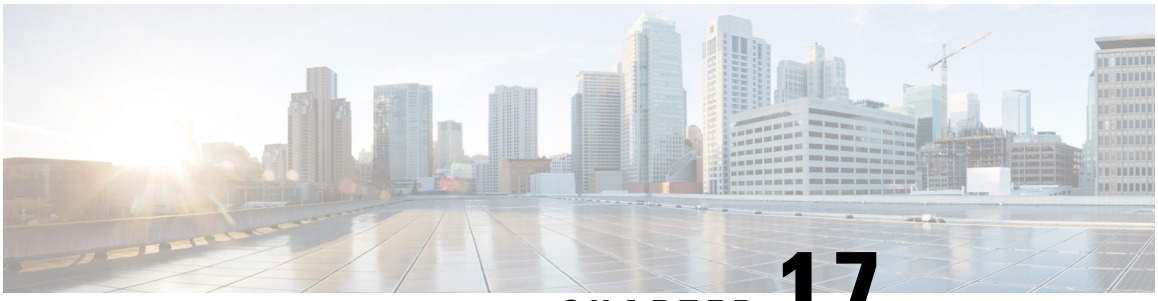

**Step 2** Click **Save**.

**Step 3** Edit the desired template(s).

**Step 4** Go to Fabrics Overview and click **Recalculate and Deploy**.

This will regenerate PTIs and the updated content will be picked up and used for the expected configuration (or intent).

**Step 5** After the contents are re-generated and deployed, uncheck the **Template In-Use Override** check box to avoid performance issues.



CHAPTER 17

# Tech Support

Initiating a Tech Support log collection attempts to query all data stores. It builds a snapshot of the current state of the system. A notification appears after the log collection is completed. You can download the log anytime.


- Log Collection, on page 449

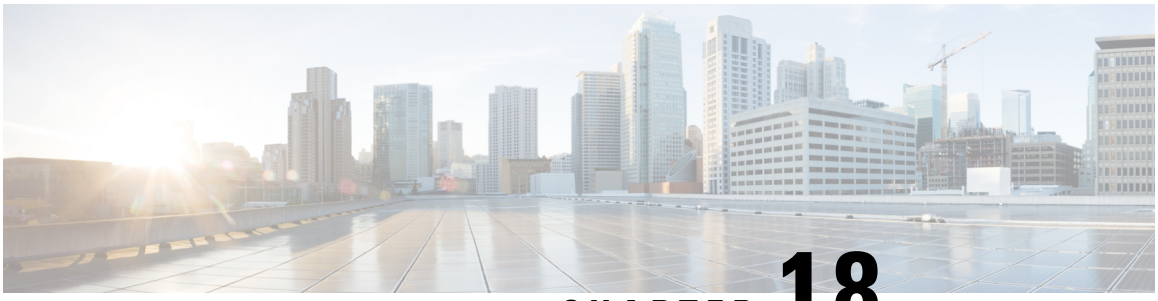

## Log Collection

Cisco Nexus Dashboard Fabric Controller allows you to collect and download logs for troubleshooting.

Click **Begin data collection** to collect logs for troubleshooting purposes.

Click **Restart log collection** to begin collecting logs. This action deletes the existing technical support logs on the server. After the collection is complete, you can download the logs for troubleshooting purposes.

Click **Download log** to download the logs to your local directory. The logs are downloaded with `.zip` extension.



CHAPTER 18

# Backup and Restore

You can take a backup manually anytime. You can also configure a scheduler to backup all fabric configurations and intents.

You can backup and restore using any of the following formats:

- **Config only**: A Config only backup is smaller. It contains the intent, dependent data, discovery information, credentials, and policies. A restore from this backup has functional fabrics, switch discovery, expected configurations, and other settings.
- **Full**: A Full backup is large. It contains current data, historical data, alarms, host information, and everything in a Config only backup. A restore from this backup has functional historical reports, metrics charts, and all base functionality.

You can restore a config-only backup or a full backup.

When restoring a backup, you can choose to do a config only restore or a full restore. A config only restore will restore only the configuration (intent, discovery information, credentials, and policies) and can be done using both config only backups and full backups. A full restore will restore the configuration and any current and historical data, charts, etc. and can be done using only full backups.

**Note** Wait for minimum of 20 minutes after fresh installation before restoring the backup data. Some applications may not be operational if the backup is restored immediately in the freshly installed setup.

### NDFC Backup and Restore Behavior in Release 12.1.2e

This table provides information about NDFC backup and restore behavior in release 12.1.2e. Anything not mentioned is assumed to be fully supported.

*Table 48: NDFC Backup and Restore Behavior*

| Feature | Config Data Backup and Restore | Operational Data Backup and Restore |
|---|---|---|
| ALL | | |
| Alarm | Supported | Not Supported |
| PM | Supported | Supported for 90 days data |

This section includes the following:


- Scheduler, on page 452
- Restore, on page 453
- Backup Now, on page 454


# Scheduler

The purpose of the scheduler is to take backups of the system, if a system needs to be restored. You must backup to a remote location.

To schedule a backups of application and configuration data from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Before you begin

If there are no scheduled backup jobs, **No Schedule set** is displayed.

### Procedure

**Step 1** Click on **No Schedule set**.

The **Scheduler** window appears.

**Step 2** Check the **Enable scheduled backups** check box.

**Step 3** Under **Type**, select your desired format to restore.

- Choose **Config only** or **Full**.

**Step 4** In the **SCP Server** field, provide the SCP server IP Address.

**Step 5** In the **File Path** field, provide the absolute path of the directory to store the backup file.

**Step 6** Enter **Username** and **Password** to the backup directory.

**Step 7** Enter the **Encryption Key** to the backup file.

You must have an Encryption Key in order to restore from the backup. The Encryption Key is used to encrypt a portion of the backup file that has sensitive information.

**Step 8** In the **Run on days** field, select the check box to schedule the backup job on one or more days.

**Step 9** In the **Start at** field, use the time picker to schedule the backup at a particular time.

The time picker is a 12-hour clock.

**Step 10** Click **Schedule backup** to run the backup job as per schedule.

# Restore

**Note** Wait for minimum of 20 minutes after fresh installation before restoring the backup data. Some applications may not be operational if the backup is restored immediately in the freshly installed setup.

From Cisco NDFC Release 12.1.2e, you can restore on a freshly installed Nexus Dashboard Fabric Controller system with no feature enabled as well as on an existing system where features have already enabled.

If the restore is done on a system with features enabled, note the following:

- You cannot restore a SAN Controller backup on a LAN Controller (and vice versa).
- You can perform only config only restore. Whether the original backup is a config only backup or a full backup, only config (non operational) data will be restored. All operational data (statistics and historical data) will be lost.

### Guidelines

You can do a restore only on a freshly installed Nexus Dashboard Fabric Controller with no features enabled.

When you migrate from L2 HA to L3 HA, check the Ignore External Service IP Configuration check box to ensure that the persistent IPs in the backup are ignored and it selects new ones during the restore. Rest of the data will be restored.

**Note** During disaster recovery, NDFC allows you to restore only on the same version on which the backup was taken.

To restore application and configuration data from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Procedure

**Step 1** Click **Restore**.

The **Restore now** window appears.

**Step 2** Under **Type**, select your desired format to restore.

- Choose **Config only** or **Full**.

**Step 3** Choose the appropriate destination where you have stored the backup file.

- Choose **Upload File** if the file is stored in a local directory.
  - **a.** Open the directory where you've saved the backup file.
  - **b.** Drag and drop the backup file to the **Restore now** window.

    or

Click Browse. Navigate to the directory where you've saved the backup file. Select the backup file and click Open.

c. Enter the **Encryption Key** to the backup file.

**Note**
You must have an Encryption Key in order to restore the backup. The Encryption Key is used to encrypt a portion of the backup file that has sensitive information.

- Choose **Import from SCP** if the backup file is stored in a remote directory.

a. In the **SCP Server** field, provide the SCP server IP Address.

b. In the **File Path** field, provide the relative file path to the backup file.

c. In the **Username** and **Password** fields, enter appropriate details.

d. In the **Encryption Key** field, enter the Encryption Key to the backup file.

**Note**
You must have an Encryption Key in order to restore the backup. The Encryption Key is used to encrypt a portion of the backup file that has sensitive information.

**Step 4** (Optional) Check the **Ignore External Service IP Configuration** check box.

If the **Ignore External Service IP Configuration** check box is selected, then the external service IP configuration is ignored. This selection allows you to take a backup on a system and restore it on a different system, with different management and/or data subnets.

This option does not have any impact during an upgrade from Cisco DCNM 11.5(x) to Cisco NDFC.

**Step 5** Click **Restore**.

The backup file appears in the table on the Backup & Restore window. The time required to restore depends on the data in the backup file.

# Backup Now

To take a backup of application and configuration data from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Click **Backup now**.

**Step 2** Under **Type**, select your desired format to restore.

- Choose **Config only** or **Full**.

**Step 3** Choose the appropriate **Destination** to store the backup file.

- Choose **Local Download** to store the backup in a local directory.
  a. Enter the **Encryption Key** to the backup file.

     **Note**
     You must have an Encryption Key in order to restore the backup. The Encryption Key is used to encrypt a portion of the backup file that has sensitive information.

  b. Click **Backup**.

     After the backup is complete, the backup file available for download from the **Backup & Restore** screen.

  c. In the Actions column, you can click on Download icon to save the backup to a local directory.

     Click on **Delete** icon to delete the backup.

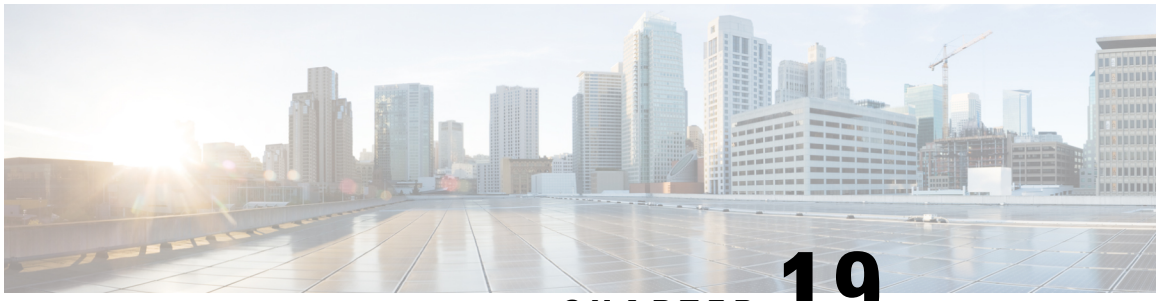
     **Note**
     You must delete the backups that are taken with **Local Download** options as soon as possible due to the limited amount of allocated disk space.

- Choose **Export to SCP** to store the backup file in a remote directory.
  a. In the **SCP Server** field, provide the SCP server IP Address.
  b. In the **File Path** field, provide the relative file path to the backup file.
  c. In the **Username** and **Password** fields, enter appropriate details.
  d. In the **Encryption Key** field, enter the Encryption Key to the backup file.

     **Note**
     You must have an Encryption Key in order to restore the backup. The Encryption Key is used to encrypt a portion of the backup file that has sensitive information.

  e. Click **Backup**.

     After the backup is complete, the backup file is saved in the remote directory.



CHAPTER 19

# NXAPI Certificates

Cisco NX-OS switches require an SSL certificate to function in NX-API HTTPS mode. You can generate the SSL certificates and get it signed by your CA. You can install the certificates manually using CLI commands on switch console or use Cisco Nexus Dashboard Fabric Controller to install these on switches.

Cisco Nexus Dashboard Fabric Controller provides a Web UI framework to upload NX-API certificates to Nexus Dashboard Fabric Controller. Later, you can install the certificates on the switches that are managed by Nexus Dashboard Fabric Controller.

**Note** This feature is supported on switches running on Cisco NXOS version 9.2(3) or higher.

• Certificate Generation and Management, on page 457

# Certificate Generation and Management

For each switch, the data center administrator generates an ASCII (base64) encoded certificate. This certificate comprises two files:

- `.key` file that contains the private key
- `.crt/.cer/.pem` file that contains the certificate

Cisco Nexus Dashboard Fabric Controller also supports a single certificate file that contains an embedded key file, that is, the `.crt/.cer/.pem` file, which can also contain the contents of the `.key` file.

Nexus Dashboard Fabric Controller doesn't support binary encoded certificates, that is, the certificates with the `.der` extension are not supported. You can protect the key file with a password for encryption. Cisco Nexus Dashboard Fabric Controller does not mandate encryption; however, as this is stored on Nexus Dashboard Fabric Controller, we recommend that you encrypt the key file. Nexus Dashboard Fabric Controller supports AES encryption.

You can either choose CA-signed certificates or self-signed certificates. Cisco Nexus Dashboard Fabric Controller does not mandate the signing; however, the security guidelines suggest you use the CA-signed certificates.

You can generate multiple certificates meant for multiple switches, to upload to Nexus Dashboard Fabric Controller. Ensure that you name the certificates appropriately, to help you choose the switch meant for that certificate.

You can upload one certificate and the corresponding key file, or bulk upload multiple certificates and key files. After the upload is complete, you can view the upload list before installing these on the switches. If a certificate file that contains an embedded key file is uploaded, Nexus Dashboard Fabric Controller derives the key automatically.

Certificate and the key file must have the same filename. For example, if a certificate filename is `mycert.pem`, the key filename must be `mycert.key`. If the certificate and key pair filenames are not the same, then Nexus Dashboard Fabric Controller will not be able to install the certificate on the switch.

Cisco Nexus Dashboard Fabric Controller allows you to bulk install the certificates to the switches. Because bulk installation uses the same password, all encrypted keys must be encrypted with the same password. If the password is different for a key, you cannot install the certificate in bulk mode. Bulk mode installation allows you to install encrypted and unencrypted keys certificates together, but all the encrypted keys must have the same password.

When you install a new certificate on the switch, it replaces the existing certificate and replaces it with the new certificate.

You can install the same certificate on multiple switches; however, you cannot use the bulk upload feature.



**Note** Nexus Dashboard Fabric Controller doesn't enforce the validity of certificates or options provided in it. It is up to you and the requirements on the switch to follow the convention. For example, if a certificate is generated for Switch-1 but it is installed on Switch-2, Nexus Dashboard Fabric Controller doesn't enforce it; switches may choose to accept or reject a certificate based on the parameters in the certificate.

### NX-API Certificate Verification by Cisco Nexus Dashboard Fabric Controller

From release 12.0.1a onwards, Cisco Nexus Dashboard Fabric Controller supports a capability to verify NX-API certificates offered by switches. The NX-API requests done by Cisco Nexus Dashboard Fabric Controller require SSL connection, and switches act like SSL server and offer server certificate as part of SSL negotiations. If provided a corresponding CA certificate, Cisco Nexus Dashboard Fabric Controller can verify it.



**Note** By default, NX-API certificate verification is not enabled because it requires all switches in the data center to have the CA-signed certificates installed, and Cisco Nexus Dashboard Fabric Controller is fed all the corresponding CA certificates.

Cisco Nexus Dashboard Fabric Controller NX-API certificate management provides two functionalities named as Switch Certificates and CA Certificates to manage the same.

# Switch Certificates

### Uploading Certificates

To upload the certificates onto Nexus Dashboard Fabric Controller, perform the following steps:

1. Click **Upload Certificate** to upload the appropriate certificate file.

2. Browse your local directory and choose the certificate key pair that you must upload to Nexus Dashboard Fabric Controller.

   You can choose certificates with extension `.cer/.crt/.pem + .key` file separately.

   Cisco Nexus Dashboard Fabric Controller also allows you to upload a single certificate file that contains an embedded key file. The key file is automatically derived after upload.

3. Click **Upload** to upload the selected files to Nexus Dashboard Fabric Controller.

   A successful upload message appears. The uploaded certificates are listed in the table.

   The table shows the Status as UPLOADED. If the certificate is uploaded without the key file, the status shows KEY_MISSING.

### Assigning Switches and Installing Certificates

To install certificates on the switches using Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Select one or multiple certificates check box.
2. From the **Actions** drop-down list, select **Assign Switch & Install**.
3. In the **NX API Certificate Credentials** field, provide the password which was used to encrypt the key while generating the certificates.

   The **Password** field is mandatory, however, if the keys were not encrypted using a password, any random string you can enter, for example, test, install, and so on. In case of unencrypted files, passwords are not used, but you still need to enter any random string because it is bulk mode.



**Note** You can install unencrypted and encrypted keys and a certificate in a single bulk install; however, you must provide the key password used for encrypted keys.

4. For each certificate, click on the **Assign** arrow and select the switch to associate with the certificate.
5. Click **Install Certificates** to install all the certificates on their respective switches.

### Unlinking and Deleting Certificates

After the certificates are installed on the switch, Nexus Dashboard Fabric Controller cannot uninstall the certificate from Nexus Dashboard Fabric Controller. However, you can always install a new certificate on the switch. The certificates that are not installed on the switches can be deleted. To delete the certificate installed on the switch, you must unlink the certificate from the switch, and then delete it from Nexus Dashboard Fabric Controller.



**Note** Unlinking the certificate from the switch does not delete the certificate on the switch. The certificate still exists on the switch. Cisco Nexus Dashboard Fabric Controller cannot delete the certificate on the Switch.

To delete certificates from Nexus Dashboard Fabric Controller repository, perform the following steps:

1. Select the certificate(s) that you need to delete.

2. From the **Actions** drop-down list, select **Unlink**.

   A confirmation message appears.

3. Click **OK** to unlink the selected certificates from the switches.

   The status column shows UPLOADED. The Switch column shows NOT_INSTALLED.

4. Select the certificate that is now unlinked from the Switch.

5. From the **Actions** drop-down list, select **Delete**.

   The certificate is deleted from Nexus Dashboard Fabric Controller.

# CA Certificates

### Uploading Certificates

To upload the certificates onto Nexus Dashboard Fabric Controller, perform the following steps:

1. On **CA Certificates** tab, click **Upload Certificate** to upload the appropriate license file.

2. Browse your local directory and choose the certificate-key pair that you must upload to Nexus Dashboard Fabric Controller.

   You can upload certificates with the `.cer/.crt/.pem` file extensions.



**Note** The CA Certificates are public certificates and do not contain any keys; also, keys are not needed for this operation. This is the certificate which Cisco Nexus Dashboard Fabric Controller must verify the NX-API certificates offered by the switches. In other words, the CA certificates are consumed by Cisco Nexus Dashboard Fabric Controller and never installed on the switches.

3. Click **Upload** to upload the selected files to Nexus Dashboard Fabric Controller.

   A successful upload message appears. The uploaded certificates are listed in the table.

### Deleting Certificates

To delete CA certificates, choose **Actions** from drop-down list, click **Delete**.

### Enabling NX-API Certificate Verification

The NX-API certificate verification is enabled using the toggle button on the CA Certificates page. However, this must be done only after all the switches managed by Cisco Nexus Dashboard Fabric Controller are installed with CA-signed certificates and the corresponding CA Root certificates (one or more) are uploaded to Cisco Nexus Dashboard Fabric Controller. When this is enabled, the Cisco Nexus Dashboard Fabric Controller SSL client starts verifying the certificates that are offered by the switches. If the verification fails, the NX-API calls will also fail.



**Note**

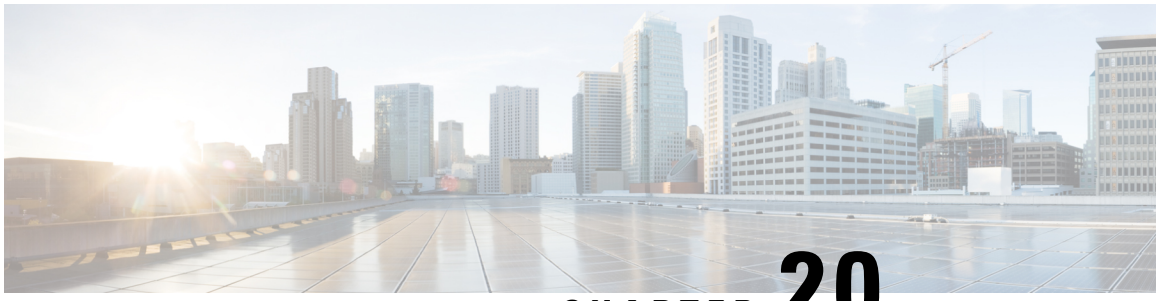
- Verification of the NX-API certificates cannot be enforced per switch; it is for either all or none. Hence, it is important that the verification is enabled only when all the switches have their corresponding CA-signed certificates installed.
- It is also required that all the CA certificates are installed on the Cisco Nexus Dashboard Fabric Controller.
- When an NX-API call fails for a given switch because of verification issues, you can use the toggle button to disable enforcement, and all goes back to the previous state without any consequences.
- Because of the above points, you must enable the enforcement during a maintenance window.



PART V

# L4-L7 Services


- L4-L7 Services Configuration, on page 465
- L4-L7 Services Use Cases, on page 489



CHAPTER 20

# L4-L7 Services Configuration

Cisco Nexus Dashboard Fabric Controller 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 L4-L7 service devices. You can add a L4-L7 service node, create route peering between the L4-L7 service node and the L4-L7 service leaf switch, and then selectively redirect traffic to these L4-L7 service nodes.

**Note** This is a preview feature in Nexus Dashboard Fabric Controller, Release 12.0.1a. We recommend that you use this feature marked as BETA in your lab setup only. Do not use these features in your production deployment.

From NDFC Release 12.0.2a, you can use this feature in your production environment.

• L4-L7 Services, on page 465

# L4-L7 Services

UI Path: **LAN** > **Services**

Alternatively, you can navigate from **LAN** > **Fabrics** > **Fabric Overview** > **Services**

Cisco provides ability to insert 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 switch, and then selectively redirect traffic to these service nodes.

You can also watch a video Service Redirection that demonstrates how to orchestrate a L4-L7 service appliance with a VXLAN Fabric in a data center managed by Cisco Nexus Dashboard Fabric Controller. This demo covers provisioning, defining of service policies, and monitoring of redirected flows.

### Service Nodes

You have to create an External Connectivity Network and specify that a service node resides in that External Connectivity Network during service node creation. Nexus Dashboard Fabric Controller does not auto-detect or discover any service node. You also have to specify the service node name, type, and form factor. The name of the service node has to be unique within a fabric. The service node is attached to a leaf, border leaf, border spine, border super spine, or a border gateway. Nexus Dashboard Fabric Controller does not define a new switch role for a service switch.

Remove the fabric from fabric Monitor Mode to display the icon to delete the service node. This delete icon will be shown only to users with admin role access.

Nexus Dashboard Fabric Controller manages the switches that are attached to a service node. Nexus Dashboard Fabric Controller also manages the interfaces of these attached switches. Ensure that the interfaces to which the service node is attached to are in trunk mode and do not belong to any interface group. The L4-L7 service will not change its mode. In case the attached switches are forming a vPC pair, the name of the attached switch is a combination of both switches.

Double-click on required service name to view the below tabs of the service node details window:

- Overview, on page 468
- Route Peering, on page 480
- Service Policy, on page 483

### VXLAN EVPN Multi-Site Support

This feature supports VXLAN EVPN Multi-Site. You can choose the VXLAN EVPN Multi-Site member fabric as attached fabric during service node creation, create a service node (for example, firewall, or load balancer), attach the service node to the switch in the selected VXLAN EVPN Multi-Site member fabric, define the route peering and service policies, and deploy relevant configurations on the selected VXLAN EVPN Multi-Site member fabric. For more information on the procedure to configure service, refer Configuring L4-L7 Services, on page 469.

### RBAC Support

The L4-L7 service supports Role-Based Access Control (RBAC) along with fabric access mode.

The admin, stager, and operator, are pre-defined roles in Nexus Dashboard Fabric Controller. The table given below lists the various operations that each role can perform.

| Service Operation | Service Node | Route Peering | Service Policy |
|---|---|---|---|
| **Create/Update/Delete/Import** | admin | admin, stager | admin, stager |
| **List/Export** | admin, stager, operator | admin, stager, operator | admin, stager, operator |
| **Attach/Detach** | NA | admin, stager | admin, stager |
| **Deploy** | NA | admin (blocked if fabric is in fabric monitor or read-only mode) | admin (blocked if fabric is in fabric monitor or read-only mode) |
| **Preview/Deployment History** | NA | admin, stager, operator | admin, stager, operator |

### PBR Support on WAN Interfaces of Border Switches

You can specify an arbitrary network, that has not been defined in the top-down configuration, as a source or destination network in the service policy. This helps in streamlining policy enforcement for north-south traffic. The Nexus Dashboard Fabric Controller UI lists out routed Layer-3 interfaces of all border switches, standalone or vPC, that have a VRF association. You can then choose the required interface that has to be associated with the defined policy. The border switches include border leaf, border spine, border super spine and border

gateway. There can be multiple interface associations. For example, multiple L3 interfaces, subinterfaces, and port-channels, can be selected for one border switch. You can also select multiple border switches for interface association. For information, see NX-OS Unicast Routing Configuration Guide.

Depending on the policy direction, the border switch and interface association for ‘any’ or arbitrary network may not be needed. For example, for a forwarding policy, the border switch and interface input or route-map association is not needed for ‘any’ or arbitrary destination network. For a reversed policy, the border switch and interface or route-map association is not needed for ‘any’ or arbitrary source network.

When the policy with ‘any’ or arbitrary network is attached, the policy related CLIs are generated and associated with the selected L3 routed interfaces of the border switches. The deployment of that policy pushes the CLIs to the selected border switches. The deployment history will include the corresponding entries and can be quickly accessed using VRF filtering. The service policy stats diagram includes the PBR stats of route maps that are associated with the selected L3 routed interfaces of the border switches.

### Static Route

The L4-L7 service pushes static routes on all VTEPs, including service leaf switches, where the VRF being referenced in the static route is attached. This expedites service node failover with static routes.

# Guidelines and Limitations for L4-L7 Services

- L4-L7 Service in Nexus Dashboard Fabric Controller does not manage or provision service nodes, such as firewall, load balancer, and Virtual Network Function.
- The L4-L7 Service feature is supported only on the VXLAN BGP EVPN fabrics with the **Data Center VXLAN EVPN** template.
- The service policies defined in this feature leverage Policy-Based Routing (PBR). Refer Nexus 9000 Series NX-OS Unicast Routing Configuration Guide for PBR related configuration, constraints, and so on.
- This feature supports Cisco Nexus 9300-EX and 9300-FX platform switches as leaf, border leaf, border spine, border super spine, and border gateway switches.
- Configurations involving intra-tenant and inter-tenant firewall for L3 networks, and one-arm Virtual Network Function and one-arm and two-arm deployed load balancer are supported.
- The existing Nexus Dashboard Fabric Controller topology view is also leveraged to display redirected flows associated with the switches that the service node is attached to, and to locate specific redirected flows.
- L4-L7 Service REST APIs are accessible via Nexus Dashboard Fabric Controller packaged REST API documentation. For more information, refer Cisco Nexus Dashboard Fabric Controller REST API Reference Guide.
- L4-L7 Services generate Kafka notifications for real-time interaction.
- Load sharing is not supported.
- This feature creates, updates, and deletes the service network, as required. Service Networks cannot be created or deleted from the **LAN > Fabrics > Networks** window.

# Types of Service Devices

The L4-L7 Service in Cisco Nexus Dashboard Fabric Controller supports any vendors service node attachments. Typical service node types that are deployed in a data center are Firewalls, Load Balancers, and other Layer-4 to Layer-7 products.

Examples of supported Firewall vendors are Cisco Systems, Palo Alto Networks, Fortinet, Check Point Software Technologies, and others.

Examples of supported Load Balancer vendors are F5 Networks, Citrix Systems, A10 Networks, and others.

Note that these example lists are meant to serve as examples and not intended to be **exhaustive** lists. The L4-L7 service attachment is generic and applies to any vendors service node.

# Overview

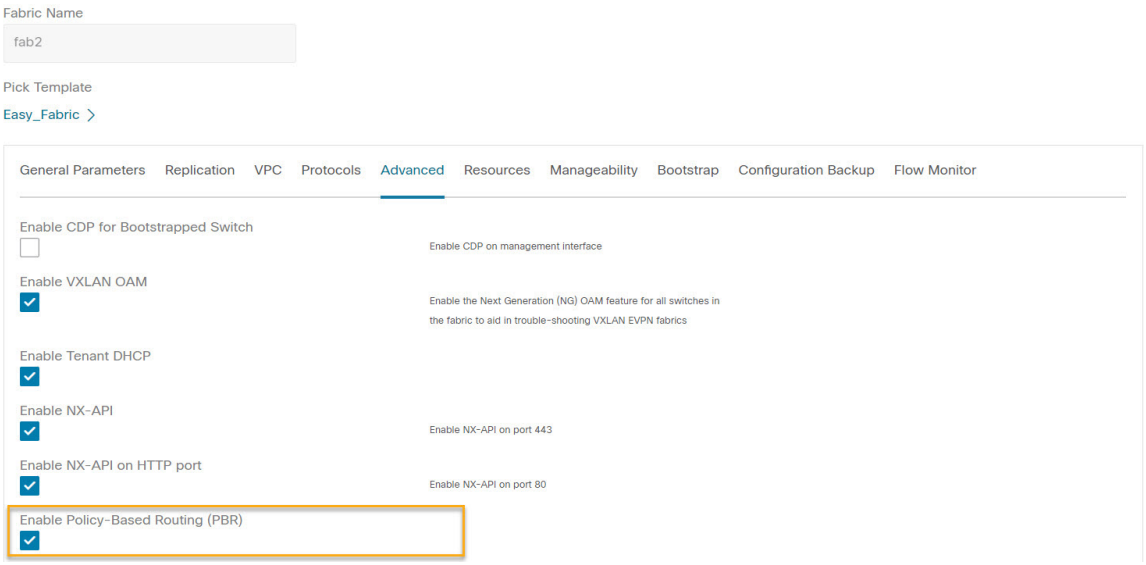
On **Overview** tab you can view **Summary**, **Route Peering**, **Service Policy** topology of selected service node.

Click **Refresh** icon to view the latest details.

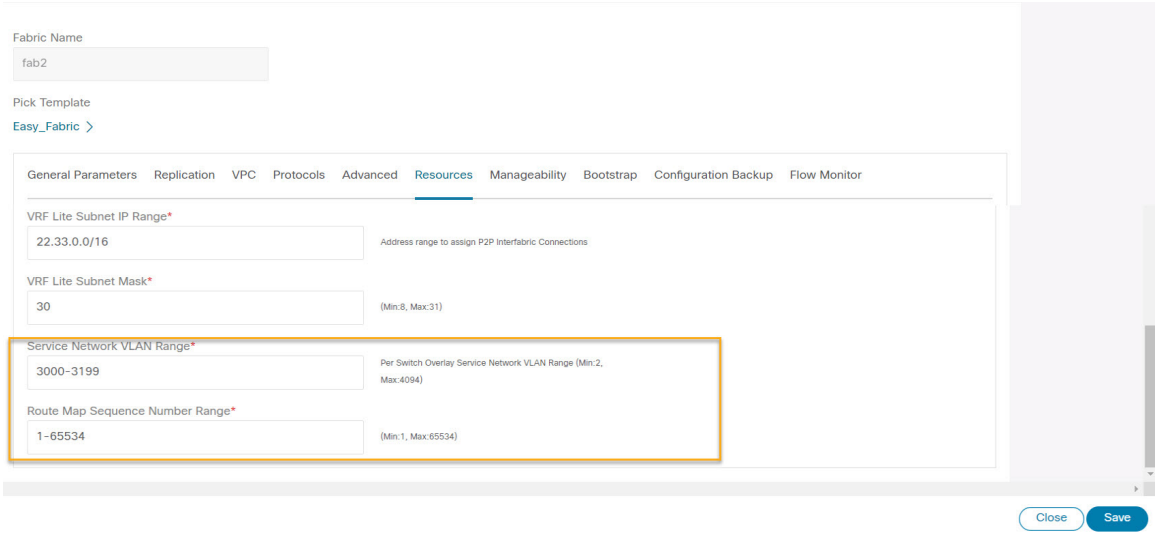
# Configuring Fabric Settings for L4-L7 Service

Certain fabric settings have to be configured to enable L4-L7 Service functionality. To configure these settings, choose **LAN > Fabrics** and then click **Actions > Create Fabric**.

The **Create Fabric** window is displayed. Provide a Fabric Name and Pick a Template. Click **Advanced**. Select the **Enable Policy-Based Routing (PBR)** checkbox to enable routing of packets based on the specified policy.




Click **Resources**. Specify a VLAN range in the **Service Network VLAN Range** field. This is a per switch overlay service network VLAN range. The minimum allowed value is 2 and the maximum allowed value is 4094. Also, specify a value for the **Route Map Sequence Number Range** field. The minimum allowed value is 1 and the maximum allowed value is 65534. Click **Save** to save the updated configuration.



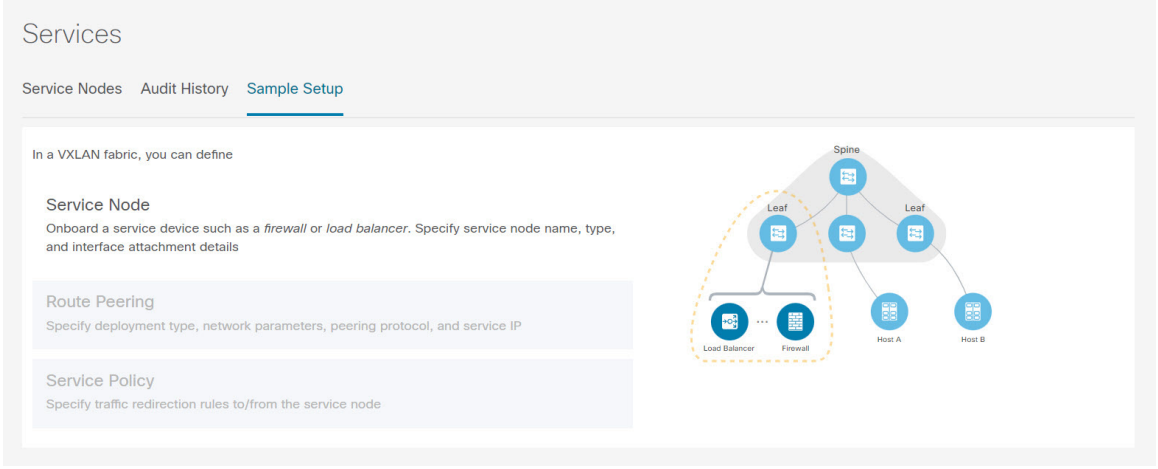

# Configuring L4-L7 Services

To launch the L4-L7 Services, or the Elastic Service, on the Cisco Nexus Dashboard Fabric Controller Web UI, choose **LAN > Services**.

You can also navigate to **Services** tab by one of the following below mentioned paths:

**LAN** > **Fabrics** > **Fabric Overview** > **Services**

**LAN** > **Switches** > **Switches Overview** > **Services**

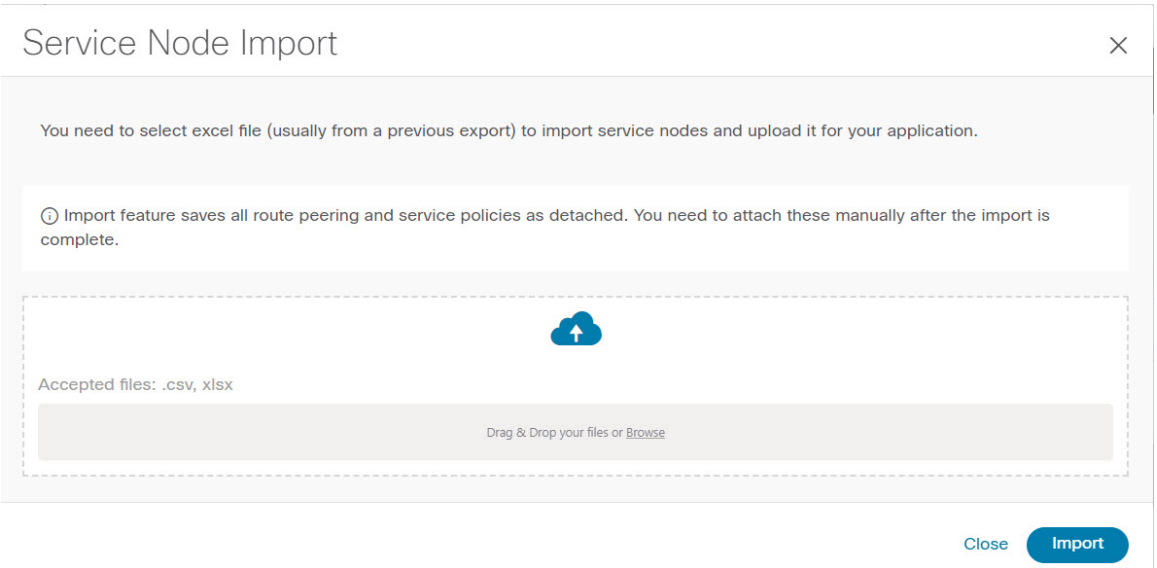



The services configuration procedure consists of the following steps:

## Adding Service Node

You can navigate to **Service Node** tab by one of the following below mentioned paths:

**LAN** > **Services**




To create service node, click **Actions** > **Add** > **Service Nodes**. The **Create New Service Node** window is displayed.

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

The **Create New Service Node** window has two sections - **Create Service Node** and **Switch Attachment**, followed by a **Link Template** drop-down list. You can select service_link_trunk, service_link_port_channel_trunk and service_link_vpc from this drop-down list based on the specified attached switch interface type.

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

### Create New Service Node

**Service Node Name**: Enter a name for the service node. The name can have alphanumeric, underscore, or dash characters.

**Service Node Type**: Select Firewall, Load Balancer, or Virtual Networking Function.

**Form Factor**: Select Physical or Virtual.

**External Fabric**: Specify the external fabric.

**Service Node Interface**: Specify the service node interface.

**Attached Fabric**: Select a fabric from the list.

**Attached Switch**: Select a switch or a switch pair from the list.

**Attached Switch Interface**: Select the interface from the list. In case the vPC pair is selected from the **Attached Leaf Switch** list, the vPC channel will be shown in the **Attached Switch Interface** list. Otherwise, the port-channel and interfaces with trunk mode are shown in the **Attached Leaf Switch Interface** list.

**Link Template**: Select the service_link_trunk, service_link_port_channel_trunk, or the service_link_vpc template. For more information on template fields, refer Templates, on page 477.

A form is displayed depending on the template used. Update all the required fields in the form and click **Save**.

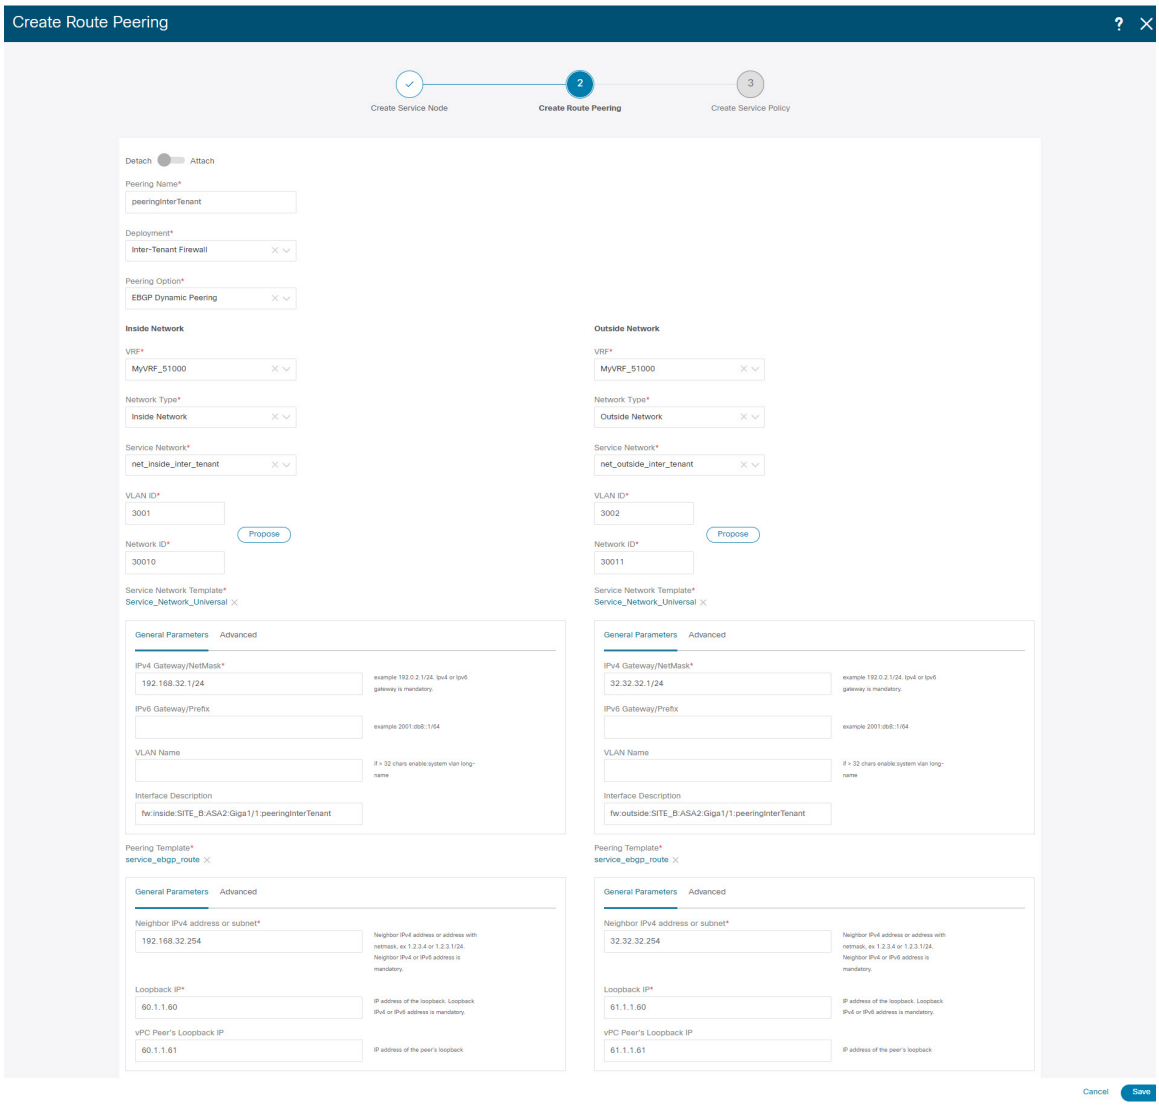
# Creating Route Peering

You can navigate to **Route Peering** tab by one of the following below mentioned paths:

**LAN** > **Services**

The fields that appear in the **Create Route Peering** window depend on the type of L4-L7 service node chosen in the **Create New Service Node** window. Depending on the type chosen (Firewall or Load Balancer or VNF), the types of deployments are Intra-Tenant Firewall, Inter-Tenant Firewall, One-Arm load balancer and Two-Arm load balancer, and One-Arm VNF.

**Note** Deletion of service network is not allowed on the **Networks** tab of detail screen from path **LAN** > **Fabrics**, click **Launch** icon **Network** window.



### Inside Network

**VRF**: Specify the VRF.

**Network Type**: Select Inside Network.

**Service Network**: Specify the name of the service network.

**VLAN ID**: Specify the VLAN ID. Valid IDs range from 2 to 3967. Click Propose to retrieve a value from the pre-defined L4-L7 service network VLAN range pool.

**Service Network Template**: Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates, on page 477.

### Outside Network

**VRF**: Specify the VRF.

**Network Type**: Select Outside Network.

**Service Network**: Specify the name of the L4-L7 service network.

**VLAN ID**: Specify the VLAN ID. Valid IDs range from 2 to 3967. Click Propose to retrieve a value from the pre-defined L4-L7 service network VLAN range pool.

**Service Network Template**: Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates, on page 477.

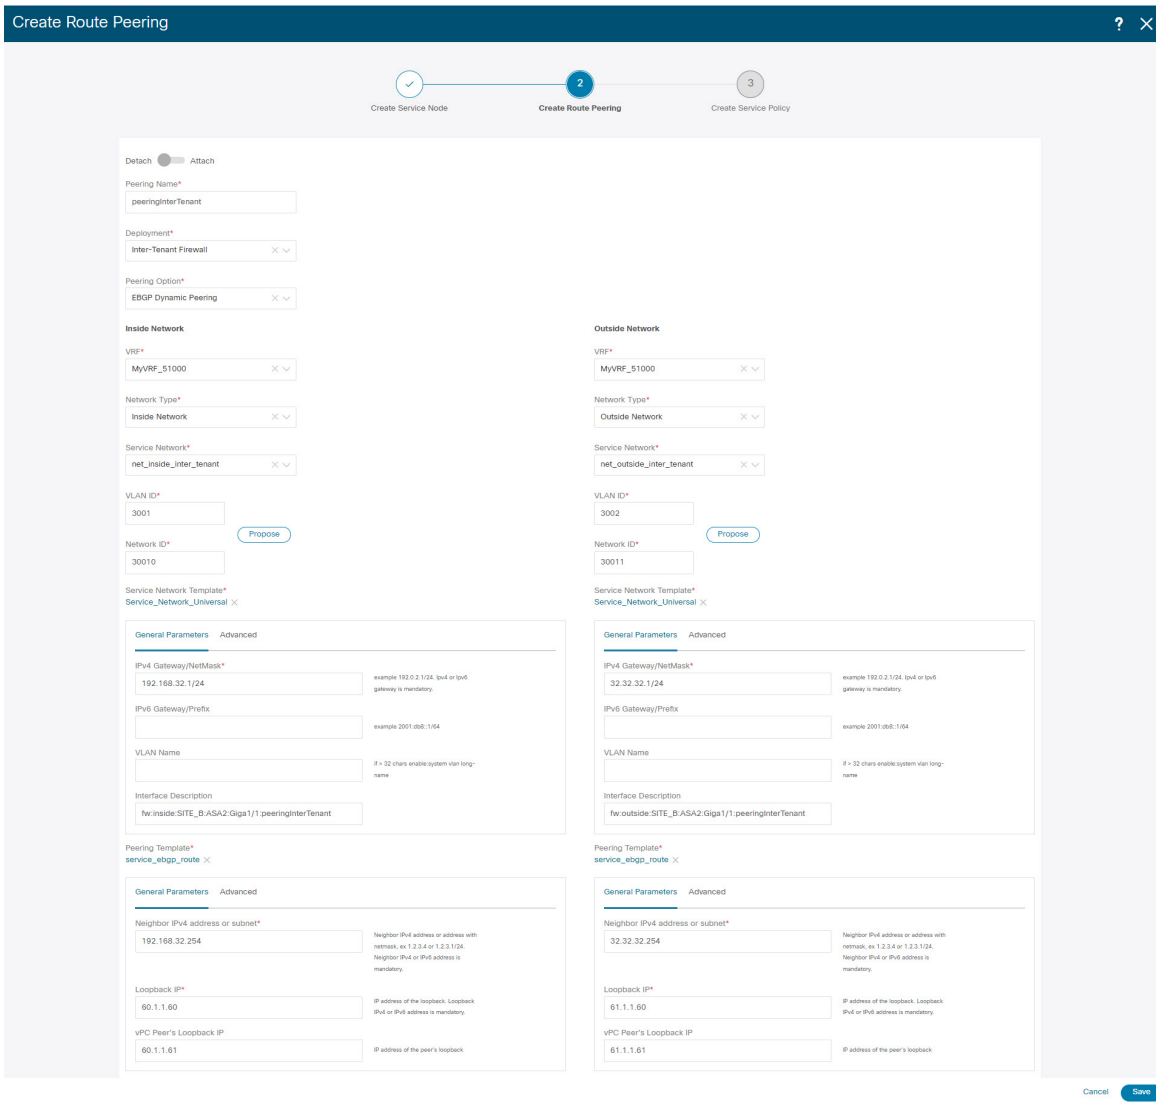
### Next Hop Section

**Next Hop IP Address**: Specify the next-hop IP address. This is the IP/VIP of the service node used for traffic redirection.

**Next Hop IP Address for Reverse Traffic**: Specify the next-hop IP address for reverse traffic. This is the IP/VIP of the service node used for traffic redirection.

### Example: Inter-Tenant Firewall Deployment

**Peering Option - Static Peering, Inside Network Peering Template - service_static_route, Outside Network Peering Template - service_static_route**




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

**Peering Name**: Specify a name for the peering. The name can have alphanumeric, underscore, or dash characters.

**Deployment**: Select Inter-Tenant Firewall.

**Peering Option**: Select Static Peering or eBGP Dynamic Peering.

### Inside Network

**VRF**: Select a VRF from the drop-down list..

**Network Type**: Select Inside Network.

**Service Network**: Provide a L4-L7 service network name.

**VLAN ID**: Specify the VLAN ID. Valid IDs range from 2 to 3967. Click **Propose** to retrieve a value from the pre-defined L4-L7 service network VLAN range pool.

**Network ID**: Specify the Network ID. Valid IDs range from to .

**Service Network Template**: Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates, on page 477.

**Peering Template** - Select service_static_route or service_ebgp_route from the drop-down list. For more information on the template fields, refer Templates, on page 477.

### Outside Network

**VRF**: Select a VRF from the drop-down list..

**Network Type**: Select Outside Network.

**Service Network**: Provide a L4-L7 service network name.

**VLAN ID**: Specify the VLAN ID. Valid IDs range from 2 to 3967. Click **Propose** to retrieve a value from the predefined L4-L7 service network VLAN range pool.

**Service Network Template**: Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates, on page 477.

**Peering Template**: Select service_static_route or service_ebgp_route from the drop-down list. For more information on the template fields, refer Templates, on page 477.

### Example: One-Arm Mode Load Balancer

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

**Peering Name**: Specify a name for the peering. The name can have alphanumeric, underscore, or dash characters.

**Deployment**: Select One-Arm F.

**Peering Option**: Select Static Peering or eBGP Dynamic Peering.

### Inside Network

**VRF**: Select a VRF from the drop-down list..

**Network Type**: Select First Mode.

**Service Network**: Provide a L4-L7 service network name.

**VLAN ID**: Specify the VLAN ID. Valid IDs range from 2 to 3967. Click **Propose** to retrieve a value from the pre-defined L4-L7 service network VLAN range pool.

**Service Network Template**: Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates, on page 477.

**Peering Template**: Select service_static_route or service_ebgp_route from the drop-down list. For more information on the template fields, refer Templates, on page 477.

**Next Hop IP Address for Reverse Traffic**: Specify the next-hop IP address for reverse traffic.

### Example: Two-Arm Mode Load Balancer

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

**Peering Name**: Specify a name for the peering. The name can have alphanumeric, underscore, or dash characters.

**Deployment**: Select Two-Arm Mode.

**Peering Option**: Select Static Peering or eBGP Dynamic Peering.

#### First Arm

**VRF**: Select a VRF from the drop-down list..

**Network Type**: Select First Arm.

**Service Network**: Provide a L4-L7 service network name.

**VLAN ID**: Specify the VLAN ID. Valid IDs range from 2 to 3967. Click **Propose** to retrieve a value from the pre-defined L4-L7 service network VLAN range pool.

**Service Network Template**: Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates, on page 477.

**Peering Template**: Select service_static_route or service_ebgp_route from the drop-down list. For more information on the template fields, refer Templates, on page 477.

#### Second Arm

**VRF**: Select a VRF from the drop-down list..

**Network Type**: Select Second Arm.

**Service Network**: Provide a L4-L7 service network name.

**VLAN ID**: Specify the VLAN ID. Valid IDs range from 2 to 3967. Click **Propose** to retrieve a value from the pre-defined L4-L7 service network VLAN range pool.

**Service Network Template**: Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates, on page 477.

#### Next Hop Section

**Next Hop IP Address for Reverse Traffic**: Specify the next-hop IP address for reverse traffic.

Now, click **Save**. The **Create Policy** window is displayed.

### Example: One-Arm Virtual Network Function

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

**Peering Name**: Specify a name for the peering. The name can have alphanumeric, underscore, or dash characters.

**Deployment**: Select One-Arm Mode.

**Peering Option**: Select Static Peering or eBGP Dynamic Peering.

**One Arm**

**VRF**: Select a VRF from the drop-down list..

**Network Type**: Select One Arm.

**Service Network**: Provide a L4-L7 service network name.

**VLAN ID**: Specify the VLAN ID. Valid IDs range from 2 to 3967. Click Propose to retrieve a value from the predefined L4-L7 service network VLAN range pool.

**Service Network Template**: Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates.

**IPv4 Gateway/Netmask**: Specify the IPv4 gateway and netmask.

**Peering Template**: Select service_static_route or service_ebgp_route from the drop-down list. For more information on the template fields, refer Templates.

**Next Hop IP Address for Reverse Traffic**: Specify the next-hop IP address for reverse traffic.

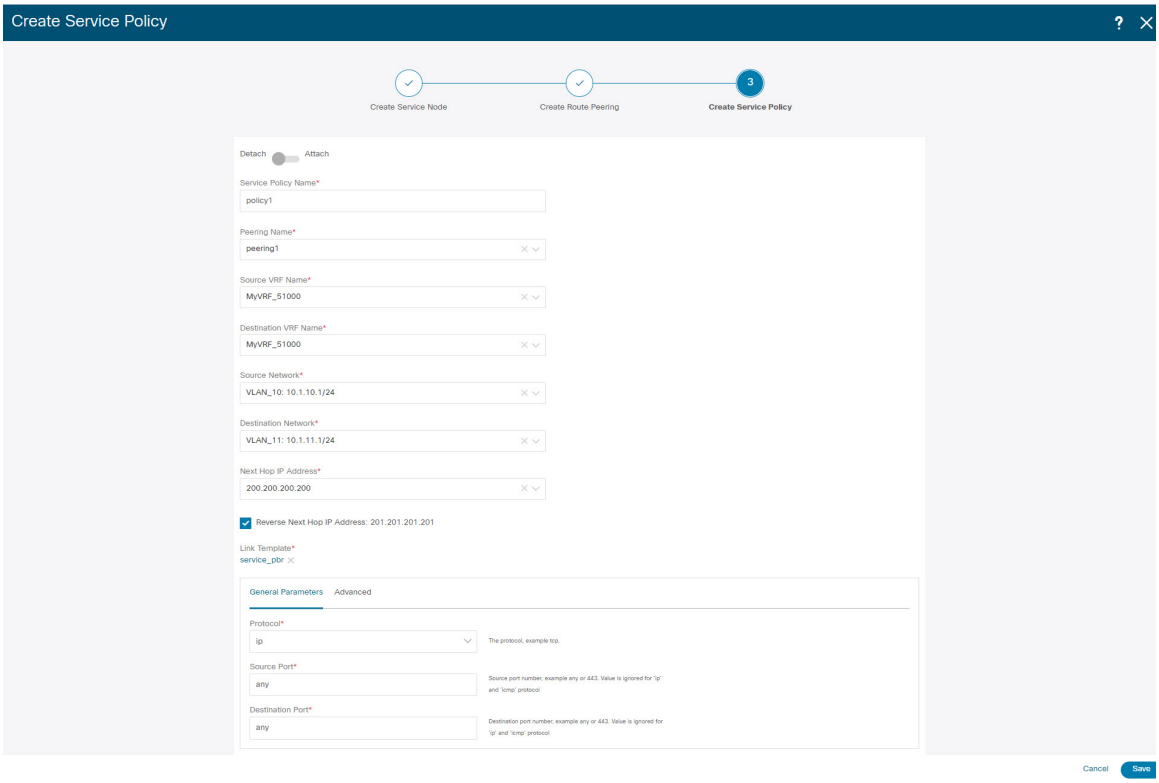
Now, click **Save**. The **Create Policy** window is displayed.

## Creating Service Policy

You can navigate to **Service Policy** tab by one of the following below mentioned paths:

**LAN** > **Services**

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



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

**Service Policy Name**: Specify a name for the policy.

**Peering Name**: Select the name of the route peering from the drop-down list.

**Source VRF Name**: Select a source VRF from the drop-down list.

**Destination VRF Name**: Select a destination VRF from the drop-down list.

**Source Network**: Select an IP address from the drop-down list.

**Destination Network**: Select a network from the drop-down list, or type in an arbitrary network with subnet info. The same is for destination network.

**Next Hop IP Address**: The next-hop IP address is displayed.

**Reverse Next Hop IP Address** - The reverse next-hop IP address is displayed. By default, the check box will be chosen.

**Link Template** : Select a template from the drop-down list. For more information on the template fields, refer Templates, on page 477.

#### General Parameters

**Protocol**: Select a protocol from the drop-down list. The options are icmp, ip, tcp, and udp.

**Source Port**: Specify a source port number. In case the ip protocol is selected, this value is ignored.

**Destination Port**: Specify a destination port number. In case the ip protocol is selected, this value is ignored.

The **Advanced** tab allow you to customize the matched traffic redirection. For example, you can specify matched traffic to be redirected using PBR, or for matched traffic to bypass a firewall and use routing table rules instead, or you can specify that any matched traffic has to be dropped. You can choose to override the route map match sequence number for prioritization. You can also customize the ACL name, however ensure that the ACL name that you specify is unique and the same name is not used for another ACL. If you do not specify the route map match sequence number or ACL name, the sequence number will be auto-populated from the designated resource pool and the ACL name will be auto-generated based on 5-tuples. For more information on the fields in the **Advanced** tab, refer Templates, on page 477.

Click **Save**. The service policy is created.



**Note** Deletion of any service network in Top-Down provisioning that is used by Services is not allowed. Deletion of any regular network that is used in a service policy is also not allowed.

# Templates

### Service Node Link Templates

#### service_link_trunk

**General Parameters tab**

**MTU**: Specifies the MTU for the interface. By default, this is set to jumbo.

**SPEED**: Specifies the speed of the interface. By default, this is set to Auto. You can change it to different supported speeds as required.

**Trunk Allowed Vlans**: Specify 'none',' all', or VLAN ranges. By default, none is specified.

**Enable BPDU Guard**: Specify an option from the drop-down list. The available options are true, false, or no. By default, no is specified.

**Enable Port Type Fast**: Check this option to enable spanning tree edge port behavior. By default, this is enabled.

**Enable Interface**: Clear the check box to disable the interface. By default, the interface is enabled.

**Advanced tab**

**Source Interface Description**: Enter a description for the source interface.

**Destination Interface Description**: Enter a description for the destination interface.

**Source Interface Freeform Config**: Enter any addition CLI for the source interface.

**Destination Interface Freeform Config**: Enter any addition CLI for the destination interface.

#### service_link_port_channel_trunk

**Port Channel Mode**: Select a port channel mode from the drop-down list. By default, active is specified.

**Enable BPDU Guard**: Specify an option from the drop-down list. The available options are true, false, or no.

**MTU**: Specifies the MTU for the interface. By default, this is set to jumbo.

**Trunk Allowed Vlans**: Specify 'none',' all', or VLAN ranges. By default, none is specified.

**Port Channel Description**: Enter a description for the port channel.

**Freeform Config**: Specify the required freeform configuration CLIs.

**Enable Port Type Fast**: Check this option to enable spanning tree edge port behavior. By default, this is enabled.

**Enable Port Channel**: Check this option to enable the port channel. By default, this is enabled.

#### service_link_vpc

This template has no specifiable parameters.

### Route Peering Service Network Template

#### Service_Network_Universal

**General Parameters tab**

**IPv4 Gateway/Netmask**: Specify the gateway IP address and mask of the service network.

**IPv6 Gateway/Prefix**: Specify the gateway IPv6 address and prefix of the service network.

**Vlan Name**: Specify a name for the VLAN.

**Interface Description**: Enter a description for the interface

**Advanced tab**

**Routing Tag**: Specify a routing tag. Valid values range from 0 to 4294967295.

### Route Peering Templates

#### service_static_route

Enter the static routes in the **Static Routes** field. You can enter one static route per line.

#### service_ebgp_route

**General Parameters tab**

**Neighbor IPv4**: Specify the IPv4 address of the neighbor.

**Loopback IP**: Specify the IP address of the loopback.

**Advanced tab**

**Neighbor IPv6**: Specify the IPv6 address of the neighbor.

**Loopback IPv6**: Specify the IPv6 address of the loopback.

**Route-Map TAG**: Specify route-map tag that is associated with the interface ID.

**Interface Description**: Enter a description for the interface.

**Local ASN**: Specify a local ASN to override the system ASN.

**Advertise Host Routes**: Select this option to enable advertisement of /32 and /128 routes to edge routers.

**Enable Interface**: Clear this option to disable the interface. By default, the interface is enabled.

### Service Policy Template

#### service_pbr

**General Parameters tab**

**Protocol**: Select a protocol from the drop-down list. The options are icmp, ip, tcp, and udp.

**Source port**: Specify a source port number. In case the ip protocol is selected, this value is ignored.

**Destination port**: Specify a destination port number. In case the ip protocol is selected, this value is ignored.

**Advanced tab**

**Route Map Action**: Select an action from the drop-down list. The options are permit or deny. If you select **permit**, the matched traffic is redirected based on the next-hop option and the defined policy. If you select **deny**, the traffic is routed based on the routing table rules.

**Next Hop Option**: Specify an option for the next-hop. The options are **none**, **drop-on-fail**, and **drop**.If you select **none**, the matched traffic is redirected based on the defined PBR rules. If you select **drop-on-fail**, the matched traffic is dropped if the specified next hop is not reachable. If you select **drop**, the matched traffic is dropped.

**ACL Name**: Specify a name for the generated access control list (ACL). If not specified, this is auto-generated.

**ACL Name for reversed traffic**: Specify a name for the ACL that is generated for reversed traffic. If not specified, this is auto-generated.

**Route map match number**: Specify a route map match number. A valid value ranges from 1 to 65535. If not specified, a route map match sequence number will be retrieved from the predefined resource pool. This number is associated with the name of the ACL.

**Route map match number for reversed traffic**: Specify a route map match number for reversed traffic. A valid value ranges from 1 to 65535. If not specified, a route map match sequence number will be retrieved from the predefined resource pool. This number is associated with the name of the ACL that has been generated for reversed traffic.

You can also customize the templates based on specific requirements.

# Route Peering

UI Path: **LAN** > **Services**, double-click on required service name to view detailed window. Navigate to **Route Peering** tab.

Alternatively, you can choose **LAN > Fabrics**, click on Fabric detail view and on the **Services** to view **Route Peering** tab.

Route peering creates service networks. Nexus Dashboard Fabric Controller supports both static route and eBGP-based dynamic route peering options. After you specify the service network and select the peering policy for the tenant, Nexus Dashboard Fabric Controller automatically creates the service network under the specified tenant. Note that the terms, tenant and VRF, will be used interchangeably in this guide.

You cannot delete the service network. Deletion of service networks is handled automatically during the service route peering deletion process. There can be multiple route peerings defined per tenant/VRF.

To create Route Peering, refer to Creating Route Peering, on page 471.

The following table describes the fields that appear on Route Peering window.

| Field | Description |
|---|---|
| **Service Network One** | |
| Peering Name | Specifies the peering name of service<br>Doube-click on **Peering Name**, detailed window appears. For more information refer to Route Peering Details, on page 482. |
| Deployment | Specifies the type of deployment. The deployment can be one of the following:<br>• Intra-tenant Firewall<br>• Inter-tenant Firewall<br>• One-arm Load Balancer |
| Peering Option | Specifies the selected peering option |
| Status | Specifies the status of service |
| Attachment Status | Specifies the status of service, whether it is attached or detached |
| VRF | Specifies the name of VRF attached with the service node |

| Field | Description |
|---|---|
| Network Name | Specifies the name of network associated with service node |
| Gateway IP | Specifies the gateway IP address of the service node |
| **Service Network Two** | |
| VRF | Specifies the name of VRF attached with the service |
| Network Name | Specifies the name of network associated with service node |
| Gateway IP | Specifies the gateway IP address |
| Next Hop IP | Specifies the hop IP address associated with the service node |
| Reverse Next Hop IP | Specifies the reverse next hop IP address associated with the service node |
| Next Hop IPv6 | Specifies the next hop IPv6 address associated with the service node |
| Reverse Next Hop IPv6 | Specifies the reverse next hop IPv6 address associated with the service node |
| Last Updated | Specifies the last modification time and date for the service node |

The following table describes the action items, in the **Actions** drop-down list, that appears in the **Route Peering** window.

| Action Item | Description |
|---|---|
| Add | Choose **Add**. The **Create Route Peering** window appears.<br>Specify the required parameters and click **Save**. |
| Edit | Choose required peering and click **Edit**. The **Edit Route Peering** window appears.<br>Use the toggle to attach or detach the route peering. When the service policy is attached or enabled, the corresponding policies are applied to the VRF (tenant), source, and destination networks.<br>Specify the required parameters and click **Save**. |
| Attach | To attach a specific route peering to a switch, choose the required peering and click **Attach**.<br>**Note**<br>Bulk attachment, detachment, preview and deployment of route peering are supported and they are limited up to 10 route-peerings only. |
| Detach | To detach a specific route peering from a switch, choose the required peering and click **Detach**. |

| Action Item | Description |
|---|---|
| Preview | To display the preview, choose the required peering and click **Preview**.<br>A **Preview Route Peering** window is displayed.<br>Select a specific switch, network, or VRF from the respective drop-down lists to display the route peerings for specific switches, networks, and VRFs. Click **Close** to close the window. |
| Deploy | To deploy a route peering, choose required peering, click **Deploy**.<br>A pop-up window appears for confirmation to deploy. Click **Deploy**. |
| Import | To import route peering information as an Excel file, click **Import**. The **Route Peering Import** window appears.<br>Click **Browse**, choose appropriate file, and then click **Import** to import information about the route peerings. |
| Export | To export route peering information as an Excel file, click **Export**. The **Route Peering Export** window appears.<br>Click **Export** to export information about the selected route peering. |
| Delete | To delete the route peering, choose appropriate route peering, and click **Delete**. |

# Route Peering Details

To view peering details window, navigate to **Services**, double-click on required service **Name**, Peering details window appears. You can view below tabs on the window:

- Overview
- Status Details
- Route Peering
- Service Policy

### Overview

The **Overivew** tab displays **Route Peering Summary** with Inside and Outside Network details, **Service Policies**, and **Service Node** as cards.

### Status Details

This tab provides a peak into the deployed configuration. Hover over the **i** icon next to the **Status Details** field in each row to display more information.

### Service Policy

Refer to Service Policy, on page 483.

### Viewing Deployment History

This tab displays deployment history of the switches and networks that are involved in the route peering. This tab displays information such as the name of the network, VRF, and switch, status, status message, status details, and time of execution.

# Service Policy

You can define service policies with any or arbitrary network and associate it with L3 routed interface on border switches. For more information, see PBR Support on WAN Interfaces of Border Switches. The L4-L7 service does not create any VRF or network other than the service networks that are defined during route peering. When you define the service policy between the created networks, the source and destination network can be a subnet, an individual IP address or the networks that are defined in the **Services** tab of fabric detail screen. Choose **LAN >Fabric**, click on Fabric detail view to view services tab. For intra-tenant firewall, one-arm and two-arm load balancer, the L4-L7 service in Nexus Dashboard Fabric Controller uses Policy-Based Routing (PBR) for service insertion. The inter-tenant firewall does not have a service policy. You only need to create a service node and route peering for inter-tenant firewall.

As the source and destination network can be attached or deployed independent of service policy deployment, the tenant/ VRF-related service policy configuration is only attached or pushed to the switch that is attached to the service node, and the source and destination network is updated with the service policy-related configuration. You can preview and confirm the generated configuration. By default, the service policy is defined but is not enabled or attached. You have to enable or attach the service policy to activate it.

The service configuration that is related to the source and destination network will be auto-processed when the source and destination networks are to be attached, or auto-updated in case the networks are already attached or deployed. By default, Nexus Dashboard Fabric Controller will collect statistics every 5 minutes and store it in the database for aggregation and analysis. By default, the statistics are stored for a maximum of 7 days.

The service insertion is effective only on the flows to be created. There is no impact on any existing flows. Deletion of a network is not allowed in case an enabled service policy is associated with that network.

The L4-L7 service integration is built on top of the easy fabric policy enforcement. Choose **LAN** > **Fabrics**, to create a VXLAN EVPN fabric and then import Cisco Nexus 9000 Series switches into the fabric with predefined fabric policies.

To create service policy, refer to Creating Service Policy, on page 476.

The following table describes the fields that appear on Route Peering window.

| Field | Description |
|---|---|
| Policy Name | Specifies the policy name of service<br>Doube-click on **Policy Name**, detailed window appears. For more information refer to Service Policy Detail section. |
| Route Peering | Specifies the route peering name |
| Status | Specifies the status of service |
| Attachment Status | Specifies the status of service, whether it is attached or detached |
| Source VRF | Specifies the name of VRF attached with the service node |

| Field | Description |
|---|---|
| Source Network | Specifies the name of source network |
| Destination VRF | Specifies the name of destination VRF attached with the service node |
| Destination Network | Specifies the name of destination network |
| Next Hop IP | Specifies the hop IP address associated with the service node |
| Reverse Next Hop IP | Specifies the reverse next hop IP address associated with the service node |
| Next Hop IPv6 | Specifies the next hop IPv6 address associated with the service node |
| Reverse Next Hop IPv6 | Specifies the reverse next hop IPv6 address associated with the service node |
| Reverse Enabled | Specifes if reverse next-hop is enabled or not. |
| Route Map Action | The options are permit or deny. If you select **permit**, the matched traffic is redirected based on the next-hop option and the defined policy. If you select **deny**, the traffic is routed based on the routing table rules. |
| Next Hop Option | Specify an option for the next-hop. The options are **none**, **drop-on-fail**, and **drop**.If you select **none**, the matched traffic is redirected based on the defined PBR rules. If you select **drop-on-fail**, the matched traffic is dropped if the specified next hop is not reachable. If you select **drop**, the matched traffic is dropped. |
| Last Updated | Displays the time at which the service policy was last updated. |

The following table describes the action items, in the **Actions** drop-down list, that appears in the **Route Peering** window.

| Action Item | Description |
|---|---|
| Add | Choose **Add**. The **Create Service Policy** window appears.<br>Specify the required parameters and click **Save**. |
| Edit | Choose required service policy and click **Edit**. The **Edit Service Policy** window appears.<br>Use the toggle to attach or detach the service policy. When the service policy is attached or enabled, the corresponding policies are applied to the VRF (tenant), source, and destination networks.<br>Specify the required parameters and click **Save**. |
| Attach | To attach a specific service policy to a switch, choose the required policy and click **Attach**.<br>**Note**<br>Bulk attachment, detachment, preview and deployment of route peering are supported and they are limited up to 10 service policy only. |

| Action Item | Description |
|---|---|
| Detach | To detach a specific service policy from a switch, choose the required service policy and click **Detach**. |
| Preview | To display the preview, choose the required peering and click **Preview**.<br>A **Preview Service Policy** window is displayed.<br>Select a specific switch, network, or VRF from the respective drop-down lists to display the service policy for specific switches, networks, and VRFs. Click **Close** to close the window. |
| Deploy | To deploy a service policy, choose required service policy, click **Deploy**.<br>A pop-up window appears for confirmation to deploy. Click **Deploy**. |
| Import | To import service policy information as an Excel file, click **Import**. The **Service Policy Import** window appears.<br>Click **Browse**, choose appropriate file, and then click **Import** to import information about the service policy. |
| Export | To export route service policy information as an Excel file, click **Export**. The **Service Policy Export** window appears.<br>Click **Export** to export information about the selected service policy. |
| Delete | To delete the service policy, choose appropriate service policy, and click **Delete**. |

## Service Policy Details

To view service policy window, navigate to **Services**, double-click on required service **Name**, service policy details window appears. You can view below tabs on the window:

- Overview
- Status Details
- Route Peering
- Service Policy

### Overview

The **Overivew** tab displays **Policy Summary**, **Service Node**, and **Route Peering** with Inside and Outside Network as cards.

### Status Details

This tab displays **Resource Type**, **Fabric Name**, **Resource Name** details associated with the selected service policy

### Statistics

This tab displays statistical information about the configured service policies. Select a time range for which the statistics should be displayed from the **Time Range** drop-down box. You can select the date from the calendar displayed on the window and the time by clicking **select time** at the bottom right corner of the window. You can also display statistics from the last 15 minutes, 1 hour, 6 hours, 1 day, 1 week, and 1 month. Select the required time range and click **Apply**. Select a switch for which the statistics should be displayed from the **Switch** drop-down list. The statistics are then displayed for the selected switch in the specified time range.

Click **Clear Stats** to reset the statistics for a specific policy on all involved switches. If multiple policies are sharing the same route map, then the statistics of other policies are also impacted.

### Viewing Deployment History

This tab displays deployment history of the switches and networks that are involved in the service policy. This tab displays information such as the name of the network, VRF, switch name, status, status message, status details, and time of execution.

# Refreshing a Service Node

To refresh the list of service node that is displayed in the **Service Nodes** window, click the **Refresh** icon .

# Viewing Audit History

To view audit history of the switches and networks that are involved in the selected service policy or route peering, click the **Audit History** tab in the **Services** window.

Audit Logs table in the Audit History window displays information about all the actions that have been performed:

| Field | Description |
|---|---|
| User Name | Specifies the user name of service node. |
| User Role | Specifies the user role name by whom latest action performed. |
| Action Taken | Specifies the latest action performed |
| Entity | Specifies the name of service node. |
| Details | Sepcifies the details of the service node |
| Status | Sepcifies the status of the service node |
| Time | Sepcifies the action time on that node |
| More Info | Click **More Info** to view detailed information of selected service node. |

- Creation of service nodes, route peering, and service policies
- Deletion of service nodes, route peering, and service policies

- Update of service nodes, route peering, and service policies
- Attachment and detachment of route peering, and service policies
- Deployment of route peering and service policies

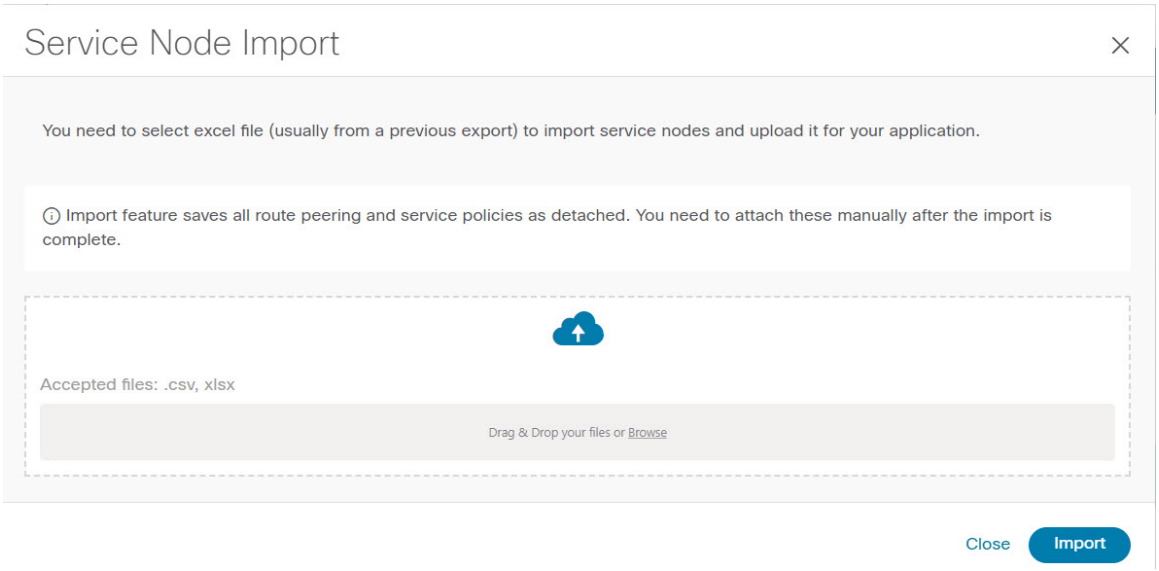
Audit logs are generated when the below actions are performed, these audit log is saved with the name of the user who has performed the action, the role of the user, the action taken, the entity on which the action was performed, details about the action, the status, and the time at which the action was performed.

To delete older audit reports, click **Action** > **Purge Audit History**, specify the maximum retained dates and confirm deletion. Note that only users with the admin role can delete audit log entries.

# Importing Service Nodes

To import service nodes from an Excel file, click **Actions > Import** on the **Service Nodes** window. The **Servie Node Import** window appears.

Click **Browser** or drag and drop your file, and click **Import** button on the **Service Node Import** window to import information about the service nodes.




You can also restore the service node level data by clicking **Actions > Import** to import data about the service nodes from an excel file.

# Exporting Service Nodes

You can back up data at the Service node level by clicking **Actions > Export** option to export data about the service nodes to an excel file. Data regarding all the service nodes, the respective route peerings, and service policy, is exported.

You can also export data for a specific Service node by selecting the node and clicking **Actions > Export**. A confirmation window appears, click **Export**.

# Editing a Service Node

To edit a service node from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose a service node from the table and click **Actions** > **Edit**.

**Step 2** The **Edit Service Node** window is displayed.

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

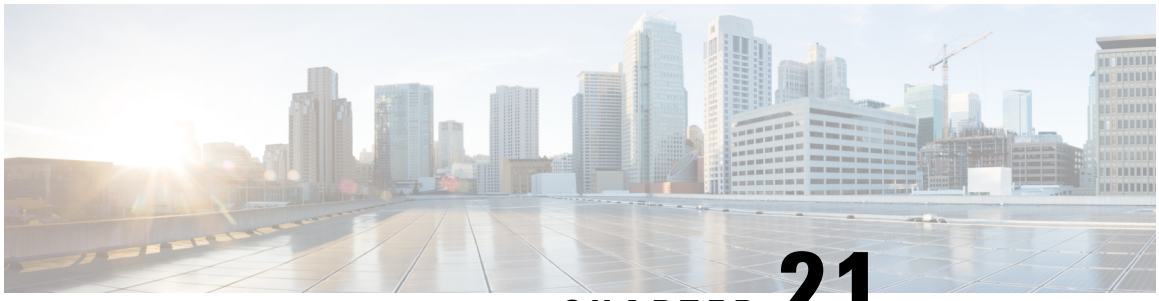
# Deleting a Service Node

To delete a service node from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

Select a service node from the table and click **Actions** > **Delete**.

**Note**
Ensure that the service node that has to be deleted has no route peering or service policies associated with it. In case there are service policies or route peering associated with the service node, the deletion is blocked with a warning indicating that any route peering or service policies associated with the service node have to be removed before deleting the service node.



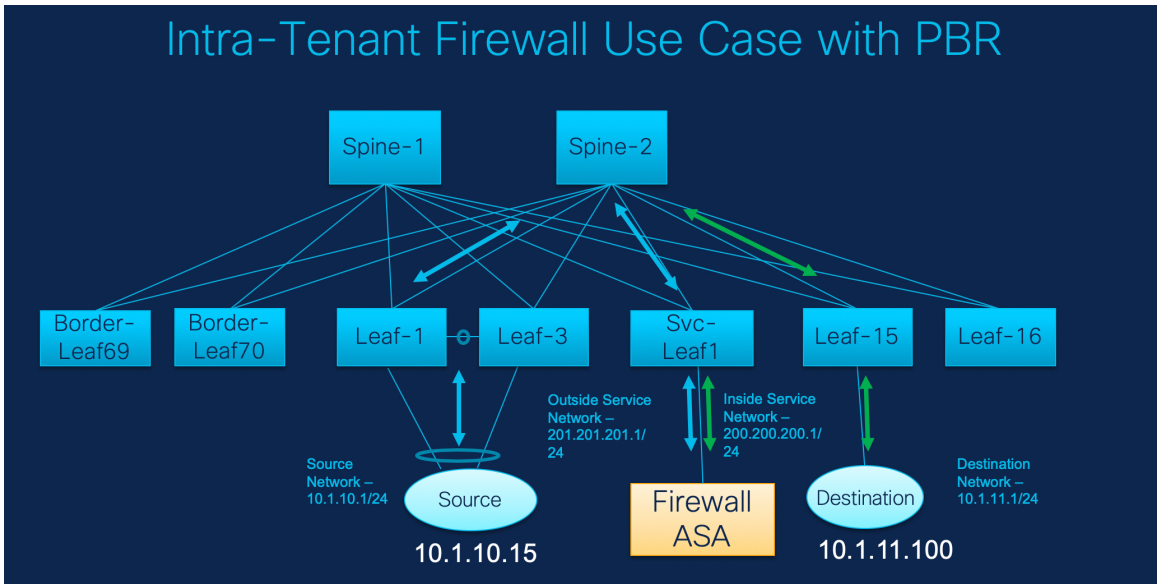
CHAPTER 21

# L4-L7 Services Use Cases


- Use Case: Intra-tenant Firewall with Policy-based Routing, on page 489
- Use Case: Inter-tenant Firewall with eBGP Peering, on page 495
- Use Case: One-arm Load Balancer, on page 500


## Use Case: Intra-tenant Firewall with Policy-based Routing

Refer the figure given below for topology details.




In this topology, Leaf1 and Leaf3 are a vPC pair and they are connected to **Source** (10.1.10.15) with the **Source Network** (10.1.10.1/24). The service leaf is connected to the virtual **Firewall ASA** and Leaf-15 is connected to **Destination** (10.1.11.100). In this use case, the source network refers to 'client' and the destination refers to 'server'.

Any traffic that is traversing from **Source** to **Destination** must go to the outside service network, and the firewall performs its function by allowing or denying traffic. This traffic is then routed to the inside service network and on to the Destination network. Since the topology is stateful, the traffic coming back from the destination to the source follows the same path.

Now, let us see how to perform service redirection in NDFC.

**Note**

- This use-case does not cover how to provision the **Site_A** VXLAN fabric. For information about this topic, refer to the Cisco Nexus Dashboard Fabric Controller for LAN Configuration Guide
- This use-case does not cover configurations on the service node (firewall or load balancer).

You can navigate to Services tab by one of the following below mentioned paths:

**LAN** > **Services**

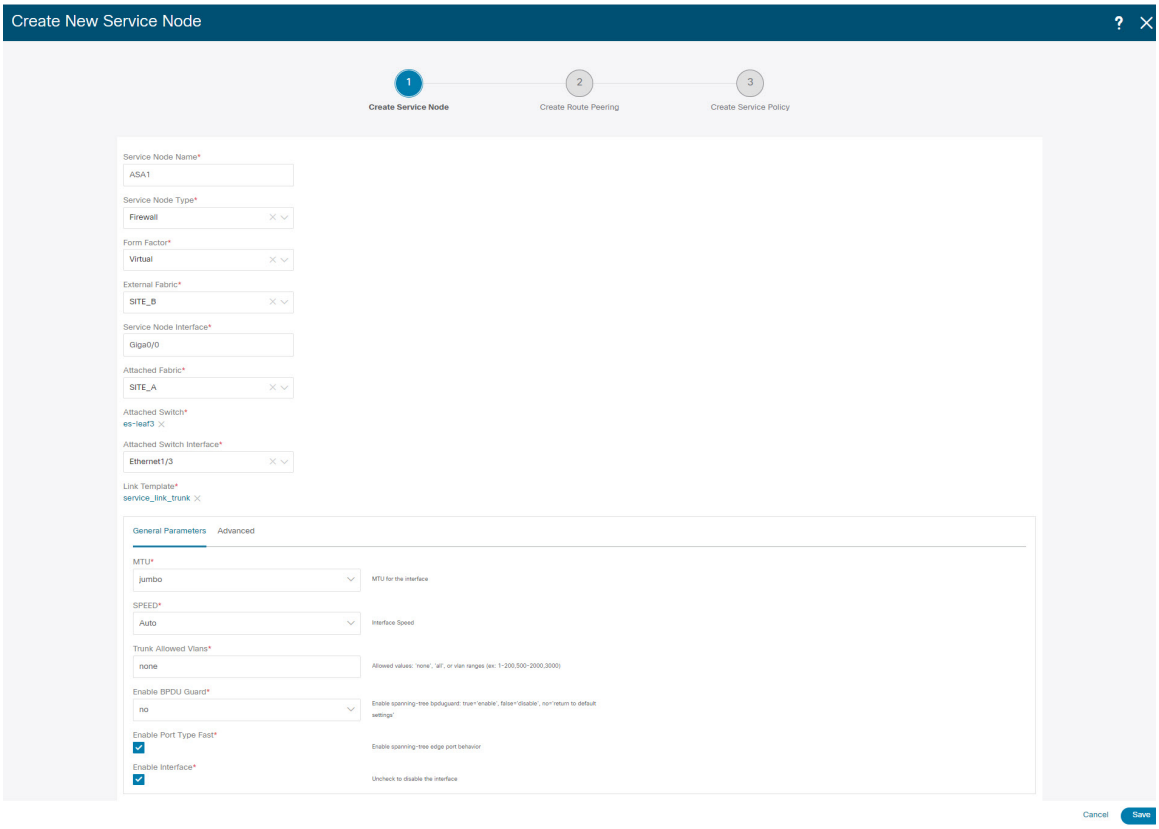
**LAN** > **Fabrics** > **Fabric Overview** > **Services**

**LAN** > **Switches** > **Switches Overview** > **Services**

# 1. Create Service Node

### Procedure

**Step 1** Navigate **LAN** > **Fabrics** > **Fabric Overview** > **Services**.




**Step 2** On **Services** tab, choose **Actions** > **Add**.

**Step 3** Enter the Service Node Name and specify **Firewall** in the **Service Node Type** dropdown box.

The **Service Node Name** must be unique.

**Step 4** From the **Form Factor** drop-down list, select **Virtual**.

**Step 5** Choose appropriate external fabric from drop-down list, select the external fabric in which the service node (for example, ASA firewall) is located.

**Note**
Ensure that the service nodes need to belong to the external fabric. This is a prerequisite before creating a service node.

**Step 6** Enter the interface name of the service node that connects to the service leaf.

**Step 7** Select the attached switch that is the service leaf, and the respective interface on the service leaf.

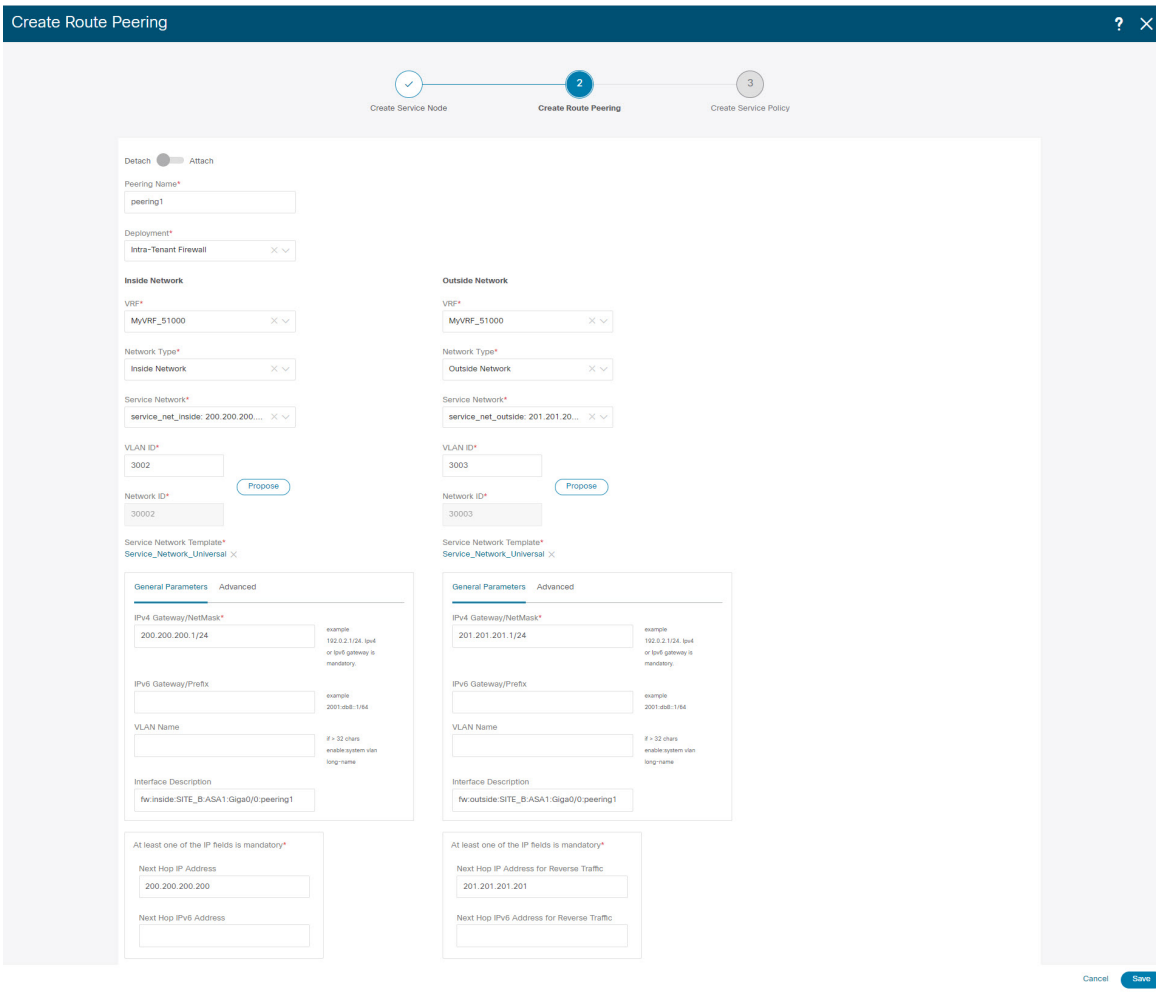
**Step 8** Choose **service_link_trunk** template. NDFC supports trunk, port channel, and vPC link templates. The available link templates in the **Link Template** drop-down list are filtered based on the selected **Attached Switch Interface** type.

**Step 9** Specify the **General Parameters** and **Advanced** parameters, if required. Some parameters are pre-filled with the default values.

**Step 10** Click **Save** to save the created service node.

# 2. Create Route Peering

Let us now configure the peering between a service leaf and a service node.




## Procedure

**Step 1** Enter the peering name and select **Intra-Tenant Firewall** from the **Deployment** drop-down list.

**Step 2** Under **Inside Network**, from the **VRF** drop-down list, select a VRF that exists and select **Inside Network** under **Network Type**.

Enter the name of the **Service Network** and specify the **Vlan ID**. You can also click **Propose** to allow NDFC to fetch the next available VLAN ID from the specified service network VLAN ID range in the fabric settings. The default **Service Network Template** is **Service_Network_Universal**.

Under the **General Parameters** tab, specify the gateway address for the service network. Specify the **Next Hop IP Address**. This next hop address has to be within the 'inside service network' subnet. Under the **Advanced** tab, the default **Routing Tag** value is 12345.

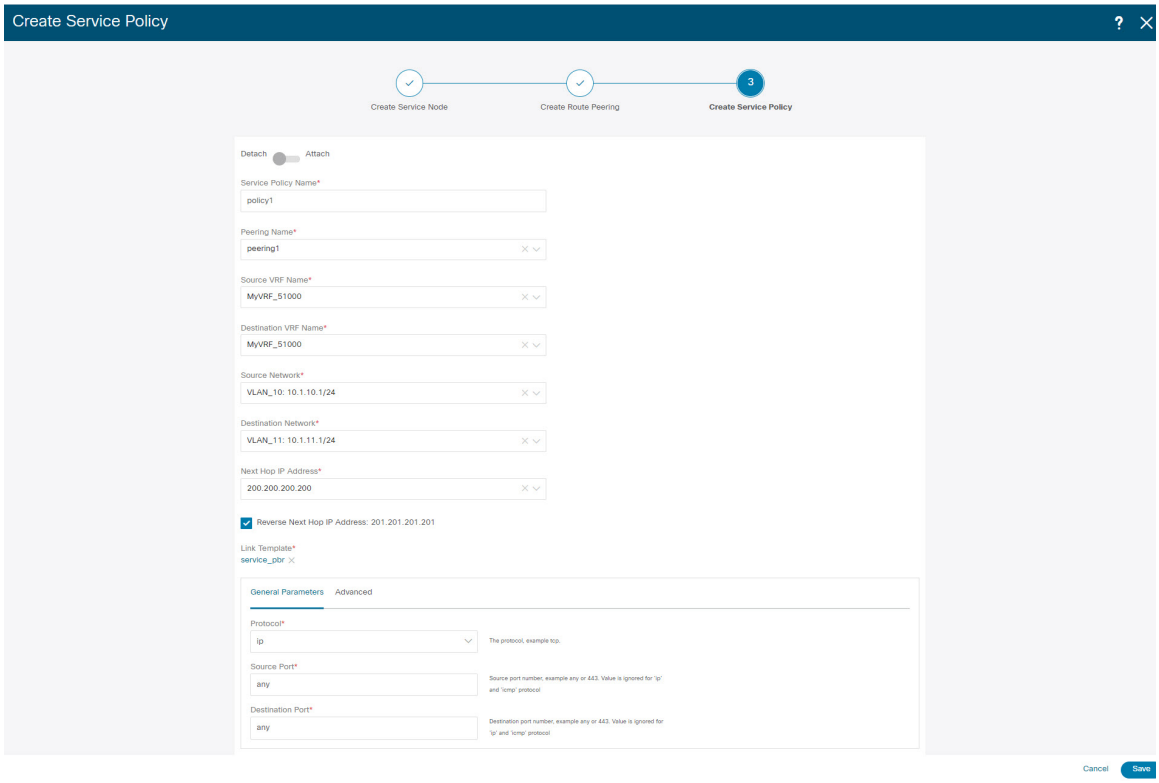
**Step 3** Specify the required parameters under **Outside Network** and specify the **Next Hop IP Address for Reverse Traffic**. This next hop address for reverse traffic needs to be within the 'outside service network' subnet.

**Step 4** Click **Save** to save the created route peering.

# 3. Create Service Policy

### Procedure

**Step 1** Specify a name for the policy and select the route peering from the **Peering Name** drop-down list.




**Step 2** Select the source and destination VRFs from **Source VRF Name** and **Destination VRF Name** drop-down lists. The source and destination VRFs for an intra-tenant firewall deployment have to be the same.

**Step 3** Select the source and destination networks from **Source Network** and **Destination Network** drop-down lists, or specify the source or destination network that is within the network subnets defined in **Fabric Overview** > **Services** window.

**Step 4** The next hop and reverse next hop fields are populated based on the values entered while creating the route peering. Select the check box next to **Reverse Next Hop IP Address** field to enable policy enforcement on reverse traffic.

**Step 5** Under the **General Parameters** tab in the policy template, select **ip** from **Protocol** dropdown list, and specify **any** in **Source Port** and **Destination Port** fields.

**Note**
For **ip** and **icmp** protocols, **any** source and destination port is used for ACL generation. You can also select a different protocol and specify the corresponding source and destination ports. NDFC converts well-known port numbers to match the format required by the switch. For example, you can convert port 80 to ‘www’.

**Step 6** Under **Advanced** tab, by default **permit** is selected for **Route Map Action** and **none** is selected for **Next Hop Option**. You can change these values, and customize the ACL name and route map match sequence number, if required. For more information, refer *Templates* section in the Layer 4-Layer 7 Service chapter.

**Step 7** Click **Save** to save the created service policy.

This completes procedures to perform and specify the flows for redirection.

# 5. Deploy Service Policy

1. On **Services** tab, on the **Service Policy** window choose the required peering.
2. Choose **Actions** > **Deploy**.

   The **Deploy Service Policy** window appears.
3. Click **Deploy** to confirm deployment.

# 4. Deploy Route Peering

1. On **Services** tab, on the **Route Peering** window choose the required peering.
2. Choose **Actions** > **Deploy**.

   The **Deploy Route Peering** window appears.
3. Click **Deploy** to confirm deployment.

# 6. View Stats

Now that the respective redirection policies are deployed, the corresponding traffic will be redirected to the firewall.

To visualize this scenario in NDFC, click the service policy, a slide-in pane appears.

You can view the cumulative statistics for a policy in a specified time range.

Statistics are displayed for:

- Forwarding traffic on the source switch
- Reversed traffic on the destination switch
- Traffic in both directions on the service switch

# 7. View Traffic Flow in Fabric Builder

The service node in the external fabric is attached to the service leaf, and this external fabric shown as a cloud icon in NDFC topology.

**Procedure**

**Step 1** Click the service leaf, a slide-in pane appears and click **Show more flows**. You can see the flows that are redirected.

**Step 2** Click **Details** in the **Service Flows** window to display attachment details.

## 8. Visualize Redirected Flows to Destination in the Topology window

**Procedure**

**Step 1** Click **Topology** and click on leafs to visualize the redirected flows to destination.

**Step 2** Select **Redirected Flows** from the drop-down list.

**Step 3** Select a policy from the drop-down list or initiate a search by entering a policy name, source network, and destination network in the search field. The search field is autopopulated based on your input.

The switches, on which the source and destination network is attached and the flows are redirected and highlighted.

**Step 4** The service node is shown as connected by a dotted line to the leaf switch on the topology window. Hover over the dotted line to get more information about the interface.

The traffic from **Source** traverses to the service leaf where the firewall is configured.

Based on firewall rules, traffic is allowed to reach the destination, Leaf 15.

# Use Case: Inter-tenant Firewall with eBGP Peering

Refer to figure given below for topology details.

In this topology, es-leaf1 and es-leaf2 are vPC border leaf switches.

Now, let us see how to perform service redirection in NDFC.

This use-case consists of the following steps:

**Note**

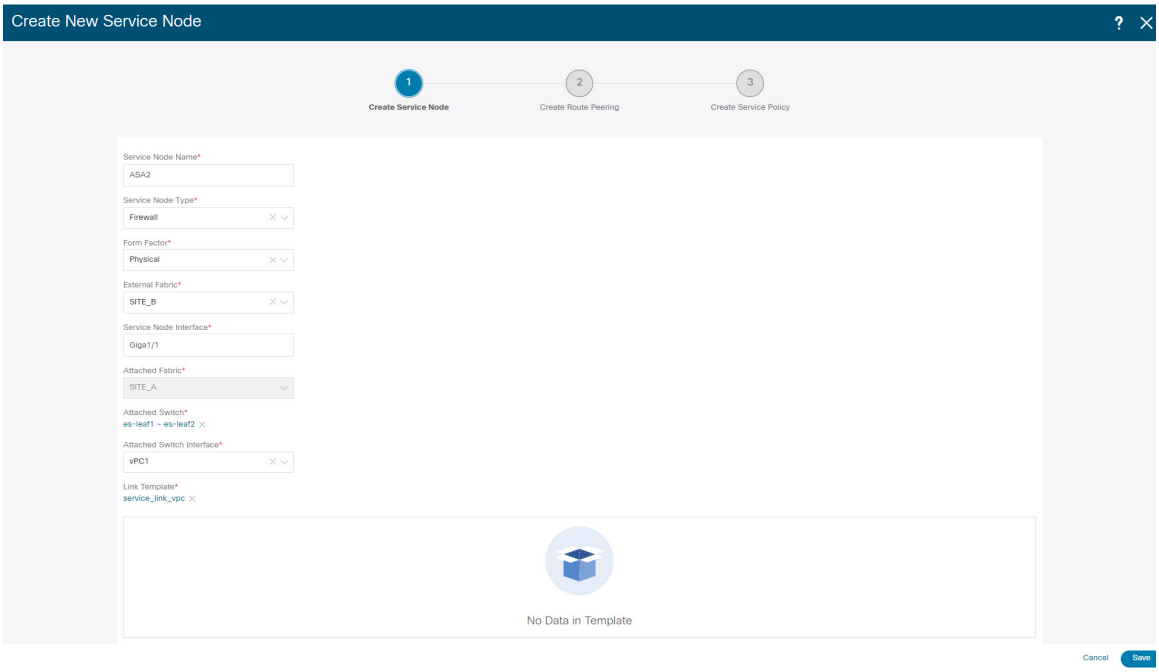
- As some steps are similar to the steps given in the Intra-tenant Firewall deployment use- case, reference links added to the steps in that use-case.
- Service policies are not applicable on Inter-tenant firewall deployments.

# 1. Create Service Node

### Procedure

**Step 1** Navigate to **LAN** > **Fabrics** > **Fabric Overview** > **Services**.

**Step 2** On **Services** tab, choose **Actions** > **Add**.




**Step 3** Enter **service node** name, choose **Firewall** in the Service Node Type dropdown box. The **Service Node Name** must be unique.

**Step 4** From the **Form Factor** drop-down list, choose **Virtual**.

**Step 5** From the **External Fabric** drop-down list, choose the external fabric in which the service node (for example, ASA firewall) is located.

**Note**

Ensure that service nodes need to belong to the external fabric. This is a prerequisite before creating a service node.

**Step 6** Enter the interface name of the service node that connects to the service leaf.

**Step 7** Select the attached switch that is the service leaf, and the respective interface on the service leaf.

**Step 8** Select the **service_link_trunk** template. NDFC supports trunk, port channel, and vPC link templates. The available link templates in the **Link Template** drop-down list are filtered based on the selected **Attached Switch Interface** type.

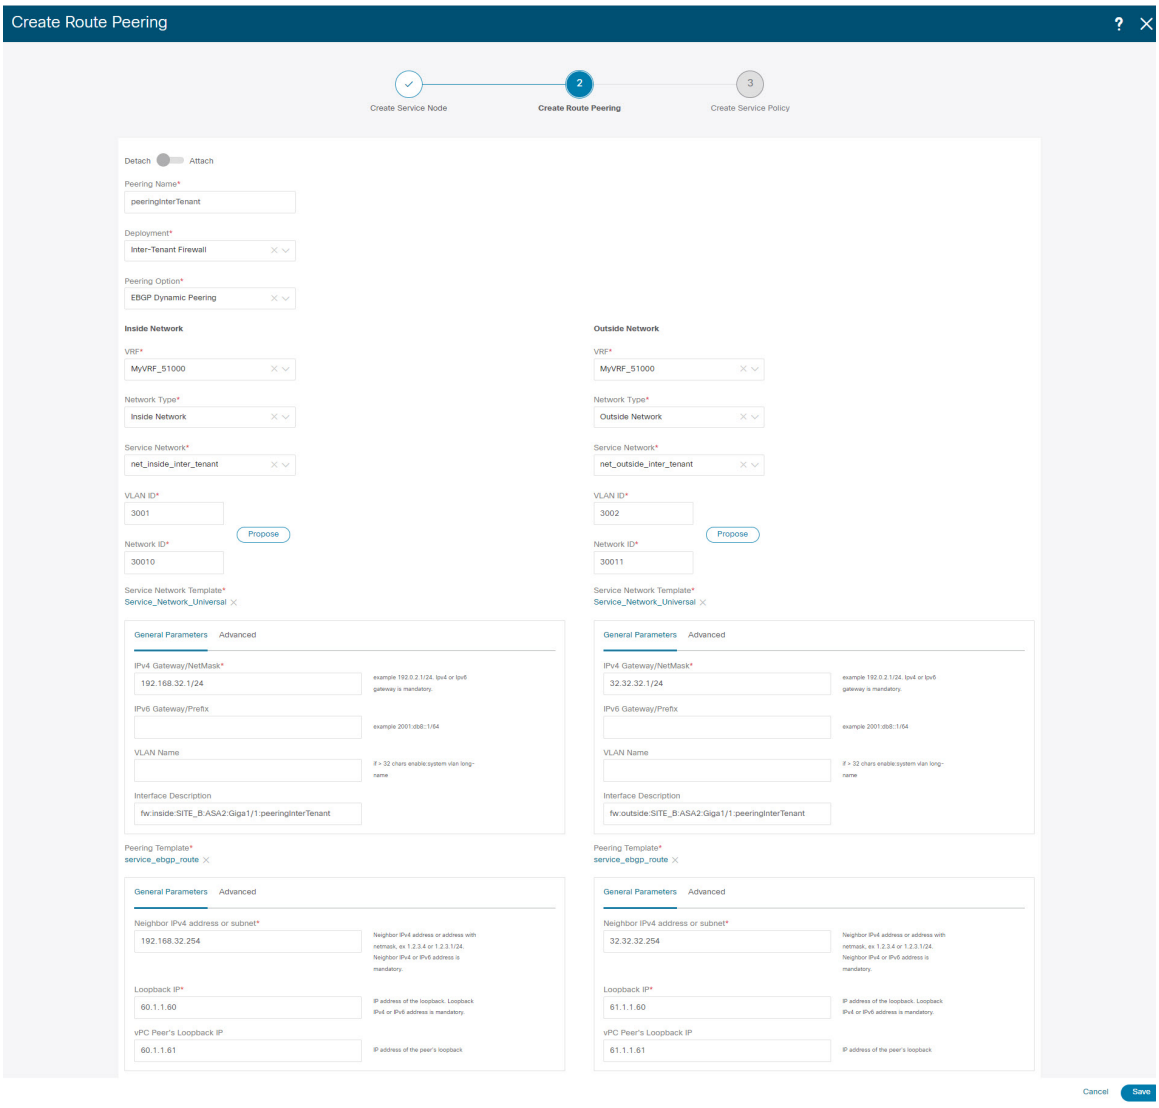
**Step 9** If required, specify **General Parameters**, and **Advanced**. Some parameters are pre-filled with default values.

**Step 10** Click **Save** to save the created service node.

**Note**
For more sample screenshots, refer 1. Create Service Node, on page 490 section in the Intra-tenant firewall with policy-based routing use case.

# 2. Create Route Peering

Let us now configure the peering between a service leaf and a service node.




## Procedure

**Step 1** Enter the peering name and select **Inter-Tenant Firewall** from the **Deployment** drop-down list. From the **Peering Option** drop-down list, select **eBGP Dynamic Peering**.

**Step 2** Under **Inside Network**from the **VRF** drop-down list, select a VRF that exists and select **Inside Network** under **Network Type**.

Enter the name of **Service Network**, specify **Vlan ID**. You can click **Propose** to allow NDFC to fetch the next available VLAN ID from specified service network VLAN ID range in the fabric settings. The default **Service Network Template** is **Service_Network_Universal**.

Under **General Parameters** tab, specify the gateway address for the service network. Specify **Next Hop IP Address**. This next hop address has to be within the ‘inside service network’ subnet. Under the **Advanced** tab, the default **Routing Tag** value is 12345.

**Step 3** The default Peering Template for eBGP dynamic peering is **service_ebgp_route**.

Under **General Parameters** tab, specify **Neighbor IPv4** address, **Loopback IP** address, and **vPC Peer's Loopback IP** address. The border switches are a vPC pair.

**Step 4** Under the **Advanced** tab, specify the **Local ASN** and select the **Advertise Host Routes** checkbox. This local ASN value is used to override the system ASN on the switch and is required to avoid routing loops.

If the **Advertise Host Routes** checkbox is selected, the /32 and /128 routes are shown. If this checkbox is not selected, the prefix routes will be shown.

By default, the **Enable Interface** checkbox is selected.

**Step 5** Specify the required parameters under **Outside Network** and specify the **Next Hop IP Address for Reverse Traffic**. This next hop address for reverse traffic needs to be within the 'outside service network' subnet.

**Step 6** The default Peering Template for eBGP dynamic peering is **service_ebgp_route**.

Under the **General Parameters** tab, **Neighbor IPv4** address, **Loopback IP** address, and **vPC Peer's Loopback IP** address. The leaf switches are a vPC pair.

**Step 7** Under the **Advanced** tab, specify the **Local ASN** and select the **Advertise Host Routes** checkbox. This local ASN value is used to override the system ASN on the switch and is required to avoid routing loops.

If the **Advertise Host Routes** checkbox is selected, the /32 and /128 routes are advertised. If this checkbox is not selected, the prefix routes will be advertised.

By default, the **Enable Interface** checkbox is selected.

**Step 8** Click **Save** to save the created route peering.

# 3. Deploy Route Peering

Refer to 4. Deploy Route Peering, on page 494 in the Intra-Tenant Firewall deployment use-case. Ensure that the **InterTenantFW** is displayed under **Deployment**.

The BGP configuration on the vPC border leaf for this use-case is given below.

```
router bgp 12345
 router-id 10.2.0.1
 address-family l2vpn evpn
  advertise-pip
 neighbor 10.2.0.4
  remote-as 12345
  update-source loopback0
  address-family l2vpn evpn
   send-community
    send-community extended
 vrf myvrf_50001
  address-family ipv4 unicast
   advertise l2vpn evpn
   redistribute direct route-map fabric-rmap-redist-subnet
   maximum-paths ibgp 2
  address-family ipv6 unicast
   advertise l2vpn evpn
   redistribute direct route-map fabric-rmap-redist-subnet
   maximum-paths ibgp 2
  neighbor 192.168.32.254
```

```
   remote-as 9876
  local-as 65501 no-prepend replace-as // Note: This configuration corresponds to the Local
 ASN template parameter value of the service_ebgp_route template of the inside network with
 VRF myvrf_50001. The no-prepend replace-as keyword is generated along with the local-as
command.
  update-source loopback2
  ebgp-multihop 5
  address-family ipv4 unicast
   send-community
   send-community extended
   route-map extcon-rmap-filter-allow-host out
vrf myvrf_50002
 address-family ipv4 unicast
  advertise l2vpn evpn
  redistribute direct route-map fabric-rmap-redist-subnet
  maximum-paths ibgp 2
 address-family ipv6 unicast
  advertise l2vpn evpn
  redistribute direct route-map fabric-rmap-redist-subnet
  maximum-paths ibgp 2
 neighbor 32.32.32.254
  remote-as 9876
  local-as 65502 no-prepend replace-as // Note: This configuration corresponds to the Local
 ASN template parameter value of the service_ebgp_route template of the outside network
with VRF myvrf_50002. The no-prepend replace-as keyword is generated along with the local-as
 command.
  update-source loopback3
  ebgp-multihop 5
  address-family ipv4 unicast
   send-community
   send-community extended
   route-map extcon-rmap-filter-allow-host out
```

The loopback interface configuration on the vPC switch es-leaf1 for this use-case is given below. The loopback interfaces in the configuration correspond to the 'Loopback IP' parameter of the **service_ebgp_route** template. Two loopback interfaces are created automatically on each vPC switch for two separate VRF instances using **Loopback IP** parameter values that are specified in the **service_ebgp_route** template.

```
interface loopback2
 vrf member myvrf_50001
 ip address 60.1.1.60/32 tag 12345
interface loopback3
 vrf member myvrf_50002
 ip address 61.1.1.60/32 tag 12345
```

The loopback interface config on vPC peer switch es-leaf2:

```
interface loopback2
 vrf member myvrf_50001
 ip address 60.1.1.61/32 tag 12345
interface loopback3
 vrf member myvrf_50002
 ip address 61.1.1.61/32 tag 12345
```

# Use Case: One-arm Load Balancer

Refer figure given below for topology details.

In this topology, es-leaf1 and es-leaf2 are vPC leafs.

Now, let us see how to perform service redirection in NDFC.

You can navigate to **Services** tab by one of the following below mentioned paths:

**LAN** > **Services**

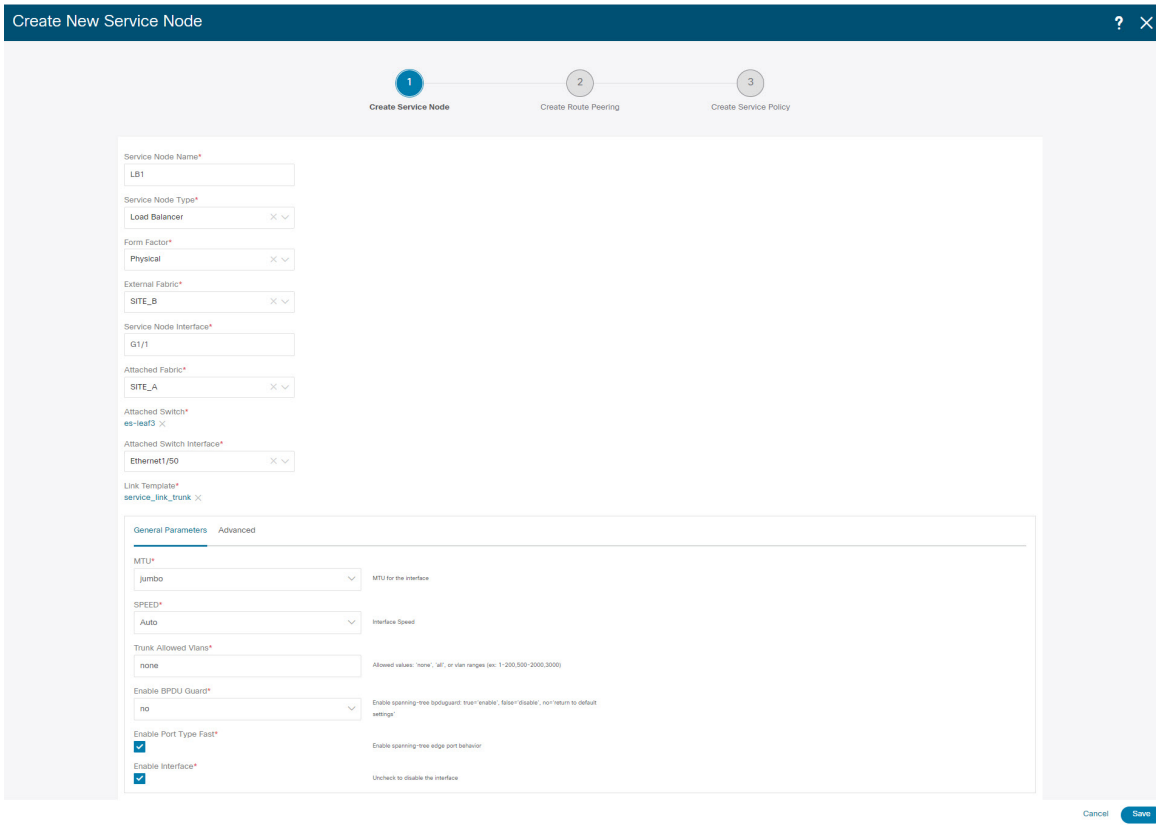
This use-case consists of the following steps:

**Note** As some steps are similar to the steps given in the Intra-tenant Firewall deployment use-case, reference links provided to the steps in that use-case.

# 1. Create Service Node

### Procedure

**Step 1** Navigate to **LAN** > **Fabrics** > **Fabric Overview** > **Services**




**Step 2** Click the **Add** icon in the **Service Nodes** window.

**Step 3** Enter the node name and specify **Load Balancer** in the **Type** dropdown box. The **Service Node Name** must be unique.

**Step 4** From the **Form Factor** drop-down list, select **Virtual**.

**Step 5** In the **Switch Attachment** section, from the **External Fabric** drop-down list, select the external fabric in which the service node (for example, ASA firewall) is located. Note that the service nodes need to belong to the external fabric. This is a prerequisite before creating a service node.

**Step 6** Enter the interface name of the service node that connects to the service leaf.

**Step 7** Select the attached switch that is the service leaf, and the respective interface on the service leaf.

**Step 8** Select the **service_link_trunk** template. NDFC supports trunk, port channel, and vPC link templates. The available link templates in the **Link Template** drop-down list are filtered based on the selected **Attached Switch Interface** type.

**Step 9** Specify **General Parameters** and **Advanced** parameters, if required. Some parameters are pre-filled with the default values.

**Step 10** Click **Save** to save the created service node.

**Note**
For more sample screenshots, refer 1. Create Service Node, on page 490 in the Intra-tenant firewall with policy-based routing use case.

# 2. Create Route Peering

Let us now configure peering between a service leaf and a service node. In this use-case, we configure static route peering.

**Procedure**

**Step 1** Enter the peering name and select **One-Arm Mode** from the **Deployment** drop-down list. Also, from the **Peering Option** dropdown list, select **Static Peering**.

**Step 2** Under **First Arm**, specify the required values. From the **VRF** dropdown list, select a VRF that exists and select **First Arm** under **Network Type**.

**Step 3** Enter the name of **Service Network** and specify **Vlan ID**. Click **Propose** to allow NDFC to fetch the next available VLAN ID from the specified service network VLAN ID range in the fabric settings. The default **Service Network Template** is **Service_Network_Universal**.

Under the **General Parameters** tab, specify the gateway address for the service network. Specify the **Next Hop IP Address**. This next hop address has to be within the first arm's subnet. Under the **Advanced** tab, the default **Routing Tag** value is 12345.

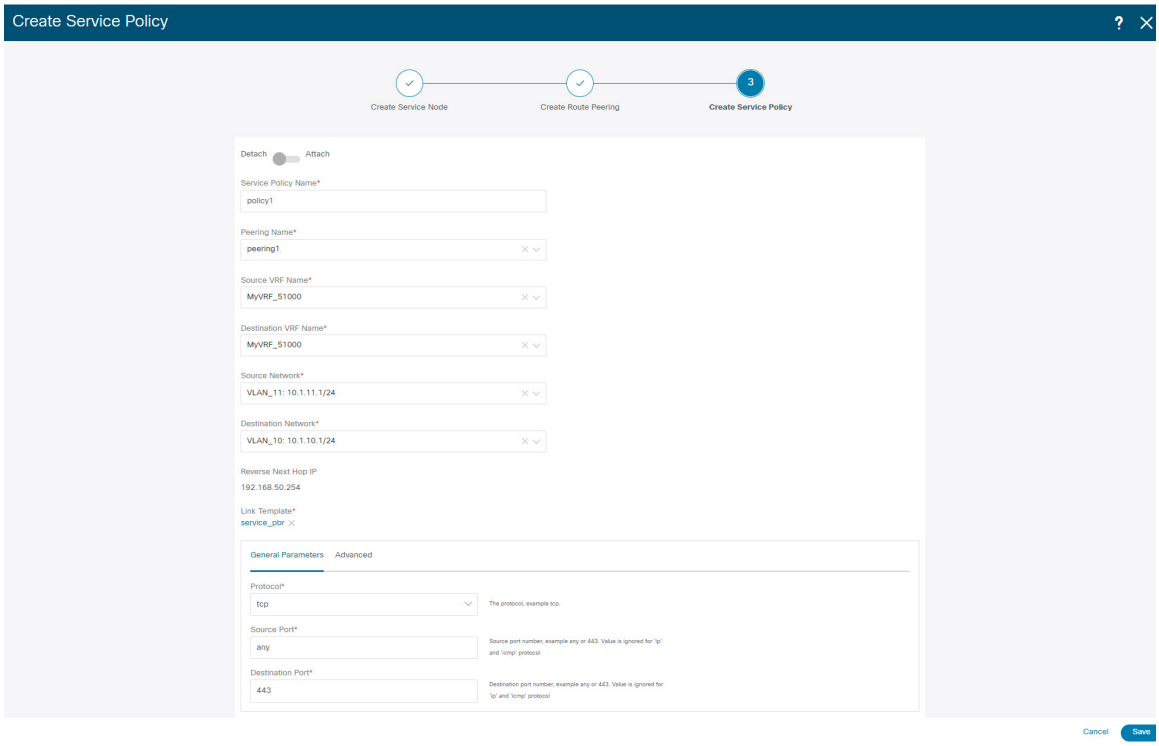
**Step 4** The default **Peering Template** is **service_static_route**. Add routes, as required, in the **Static Routes** field.

**Step 5** Specify **Next Hop IP Address** for Reverse Traffic.

**Step 6** Click **Save** to save the created route peering.

# 3. Create Service Policy

Refer to 3. Create Service Policy, on page 493 in the Intra-Tenant Firewall deployment use-case.




# 4. Deploy Route Peering

Refer to 4. Deploy Route Peering, on page 494 in the Intra-tenant Firewall deployment use-case. Note that **OneArmADC** is displayed under **Deployment**.

# 5. Deploy Service Policy

Refer to 5. Deploy Service Policy, on page 494 in the Intra-tenant Firewall deployment use-case. However, as there are two servers in this load balancer use-case, two service policies to be defined with each server network.

# 6. View Stats

Refer to 6. View Stats, on page 494 in the Intra-Tenant Firewall deployment use-case.

# 7. View Traffic Flow in Fabric Builder

Refer to 7. View Traffic Flow in Fabric Builder, on page 494 in the Intra-Tenant Firewall deployment use-case.

# 8. Visualize Redirected Flows to Destination in the Topology window

Refer to 8. Visualize Redirected Flows to Destination in the Topology window, on page 495 in the Intra-Tenant Firewall deployment use-case.

The VRF configuration on the service leaf is as given below.

```
interface Vlan2000
 vrf member myvrf_50001
 ip policy route-map rm_myvrf_50001

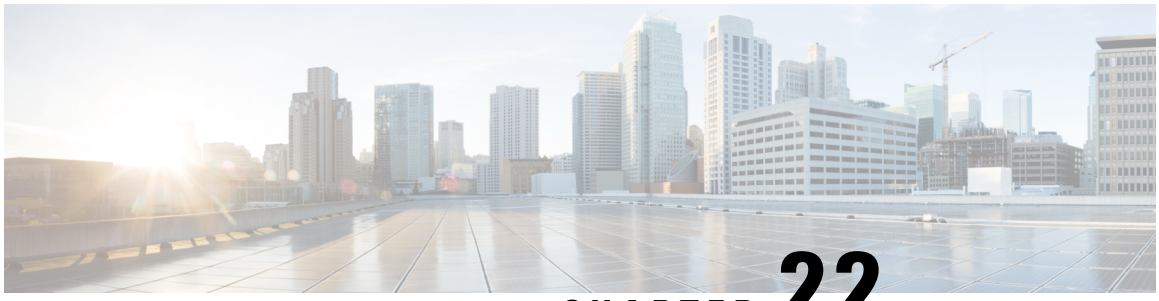
interface Vlan2306
 vrf member myvrf_50001
vrf context myvrf_50001
vni 50001
 ip route 55.55.55.55/32 192.168.50.254 // Note: This is the static route
 rd auto
 address-family ipv4 unicast
  route-target both auto
  route-target both auto evpn
 address-family ipv6 unicast
  route-target both auto
  route-target both auto evpn
router bgp 12345
 vrf myvrf_50001
  address-family ipv4 unicast
   advertise l2vpn evpn
   redistribute direct route-map fabric-rmap-redist-subnet
   redistribute static route-map fabric-rmap-redist-static
   maximum-paths ibgp 2
  address-family ipv6 unicast
   advertise l2vpn evpn
   redistribute direct route-map fabric-rmap-redist-subnet
   redistribute static route-map fabric-rmap-redist-static
   maximum-paths ibgp 2
```



# PART VI

# Service Integration


- Endpoint Locator, on page 507



CHAPTER 22

# Endpoint Locator


- Endpoint Locator , on page 507
- Monitoring Endpoint Locator, on page 521
- Disabling Endpoint Locator, on page 522


# Endpoint Locator

The Endpoint Locator (EPL) feature allows real-time tracking of endpoints within a data center. The tracking includes tracing the network life history of an endpoint and getting insights into the trends that are associated with endpoint additions, removals, moves, and so on. An endpoint is anything with at least one IP address (IPv4 and\or IPv6) and MAC address. EPL feature is also capable of displaying MAC-Only endpoints. By default, MAC-Only endpoints are not displayed. An endpoint can be a virtual machine (VM), container, bare-metal server, service appliance and so on.

**Note**

- EPL is supported for VXLAN BGP EVPN fabric deployments only in the Nexus Dashboard Fabric Controller LAN fabric installation mode. The VXLAN BGP EVPN fabric can be deployed as Easy fabric, Easy eBGP fabric, or an External fabric (managed or monitored mode). EPL is not supported for 3-tier access-aggregation-core based network deployments.
- EPL displays endpoints that have at least one IP address (IPv4 and/or IPv6). EPL is also capable of displaying MAC-Only endpoints. Select the **Process MAC-Only Advertisements** checkbox while configuring EPL to enable processing of EVPN Route-type 2 advertisements having a MAC address only. L2VNI:MAC is the unique endpoint identifier for all such endpoints. EPL can now track endpoints in Layer-2 only network deployments where the Layer-3 gateway is on a firewall, load-balancer, or other such nodes.

EPL relies on BGP updates to track endpoint information. Hence, typically the Nexus Dashboard Fabric Controller must peer with the BGP Route-Reflector (RR) to get these updates. For this purpose, IP reachability from the Nexus Dashboard Fabric Controller to the RR is required. This can be achieved over in-band network connection to the Nexus Dashboard Fabric Controller Data Network interface. There is no option to configure static routes for pods on ND, so the selected RRs must be reachable through the default data network gateway.

Some key highlights of the Endpoint Locator are:

- Support for dual-homed and dual-stacked (IPv4 + IPv6) endpoints

- Support for up to two BGP Route Reflectors or Route Servers
- Support real-time and historical search for all endpoints across various search filters such as VRF, Network, Layer-2 VNI, Layer-3 VNI, Switch, IP, MAC, port, VLAN, and so on.
- Support for real-time and historical dashboards for insights such as endpoint lifetime, network, endpoint, VRF daily views, and operational heat map.
- Support for iBGP and eBGP based VXLAN EVPN fabrics. The fabrics may be created as Easy Fabrics or External Fabrics. EPL can be enabled with an option to automatically configure the spine or RRs with the appropriate BGP configuration.
- You can enable the EPL feature for upto 4 fabrics.
- EPL is supported on Multi-Site Domain (MSD).
- IPv6 underlay is not supported.
- Support for high availability
- Support for endpoint data that is stored for up to 60 days, amounting to a maximum of 2 GB storage space.
- Support for optional flush of the endpoint data to start afresh.
- Supported scale: Maximum of 50K unique endpoints per fabric. A maximum of 4 fabrics is supported. However, the maximum total number of endpoints across all fabrics should not exceed 100K.

  If the total number of endpoints across all fabrics exceeds 100K, an alarm is generated and is listed under the **Alarms** icon at the top right of the window. This icon starts flashing whenever a new alarm is generated.
- From NDFC Release 12.0.1a, Persistent or External IP addresses are required to enable EPL. For each VXLAN fabric, a specific container is spawned running a BGP instance to peer with the spines of the fabric. This container must have a persistent IP associated that is then configured as a iBGP neighbor on the spines. A different container is used for each fabric, so the number of fabrics that are managed by NDFC where EPL is enabled decides how many persistent IP addresses must be distributed for EPL. Also, the EPL establishes iBGP sessions only over the Cisco Nexus Dashboard Data interface.
- For Virtual Cisco Nexus Dashboard deployments, enable or accept promiscuous mode on the port-groups that are associated with the Nexus Dashboard Management and/or Data vNICs where IP stickiness is required. The Persistent IP addresses are given to the pods (for example, SNMP Trap/Syslog receiver, Endpoint Locator instance per Fabric, SAN Insights receiver, and so on). Every POD in Kubernetes can have multiple virtual interfaces. Specifically for IP stickiness an extra virtual interface is associated with the POD that is allocated an appropriate free IP from the external service IP pool. The vNIC has its own unique MAC address that is different from the MAC addresses associated with the vND virtual vNICs. Moreover, all communication to and from the PODs toward an external switch goes out of the same bond interface for North-to-South traffic flows. The EPL container uses Nexus Dashboard Data Interface. The Data vNIC maps to bond0 (also known as bond0br) interface. By default, the VMware system checks if the traffic flows out of a particular vNIC is matched with the Source-MAC that is associated with the vNIC. In case of NDFC, the traffic flows are sourced with the Persistent IP addresses of the given PODs. Therefore, we must enable the required settings on the VMware side.

If you are using a Virtual Cisco Nexus Dashboard Cluster before you begin, ensure that the Persistent IP addresses, EPL feature, and required settings are enabled. See below links:

Cisco Nexus Dashboard Fabric Controller Deployment Guide

Cisco Nexus Dashboard Fabric Controller Installation and Upgrade Guide

### Backup and Restore

EPL only backups data for fabrics that EPL has been configured. If EPL is disabled for a fabric(even if EPL has previously been configured there), then you cannot backup the data for that fabric. Also, you can backup only historical data (data on the Endpoint Search page).

If a backup is initiated when EPL is enabled, then when restoring the backup, the same external data IPs that EPL was using must be available on ND. If those IPs are not available, then select the **Ignore External Service IP Configuration** option in the restore backup form. However, there are chances that the EPL pods will be brought up with different IPs, so any existing EPL policies become invalid. If EPL was previously configured with the **Configure My Fabric** option, you need to disable and enable EPL so that the old policy is cleaned up and an updated policy is deployed. If you did not use the **Configure My Fabric** option, then manually update their config with the new IPs.

# EPL Connectivity Options

Sample topologies for the various EPL connectivity options are as given below.

### NDFC Cluster Mode: Physical Server to VM Mapping

Refer to Cisco Nexus Dashboard Fabric Controller Verified Scalability Guide for more information.

# Configuring Endpoint Locator

The Nexus Dashboard Fabric Controller OVA or the ISO installation comes with two interfaces:

- Management
- Data

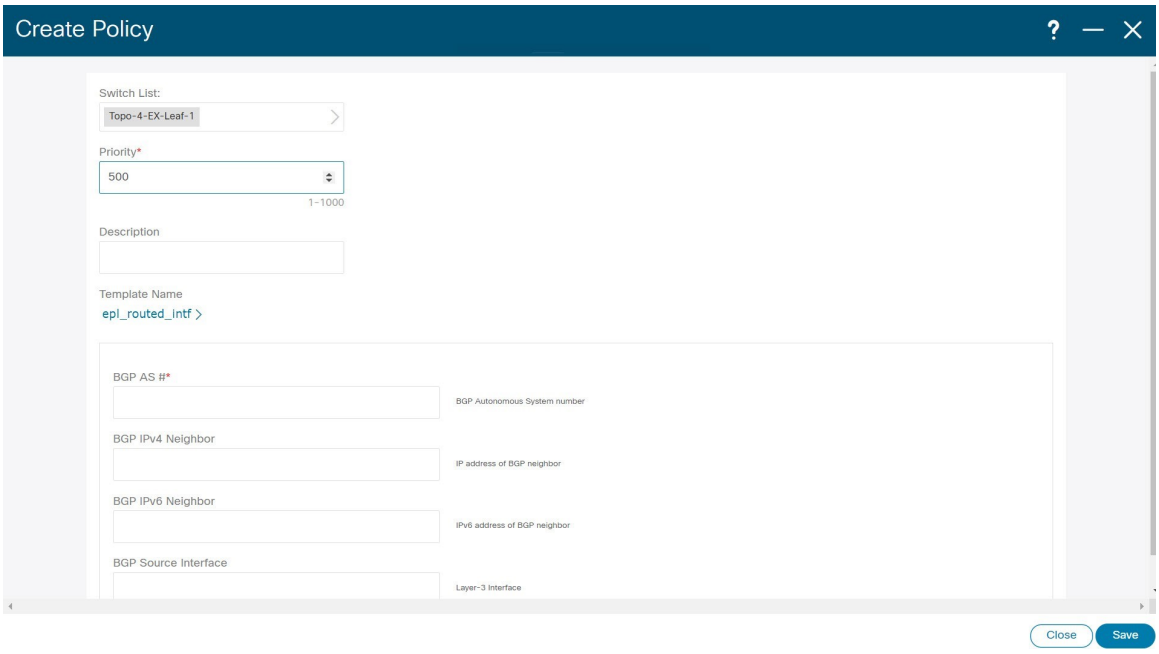
(Out-of-band or OOO) connectivity of switches via switch mgmt0 interface can be through data or Management interface. For more information refer to NDFC Installation and Upgrade Guide.

The Management interface provides reachability to the devices via the mgmt0 interface either Layer-2 or Layer-3 adjacent. This allows Nexus Dashboard Fabric Controller to manage and monitor these devices including POAP. EPL requires BGP peering between the Nexus Dashboard Fabric Controller and the Route-Reflector. Since the BGP process on Nexus devices typically runs on the default VRF, in-band IP connectivity from the Nexus Dashboard Fabric Controller to the fabric is required. The data network interface can be configured during Nexus Dashboard installation. You can't modify the configured in-band network configurations.

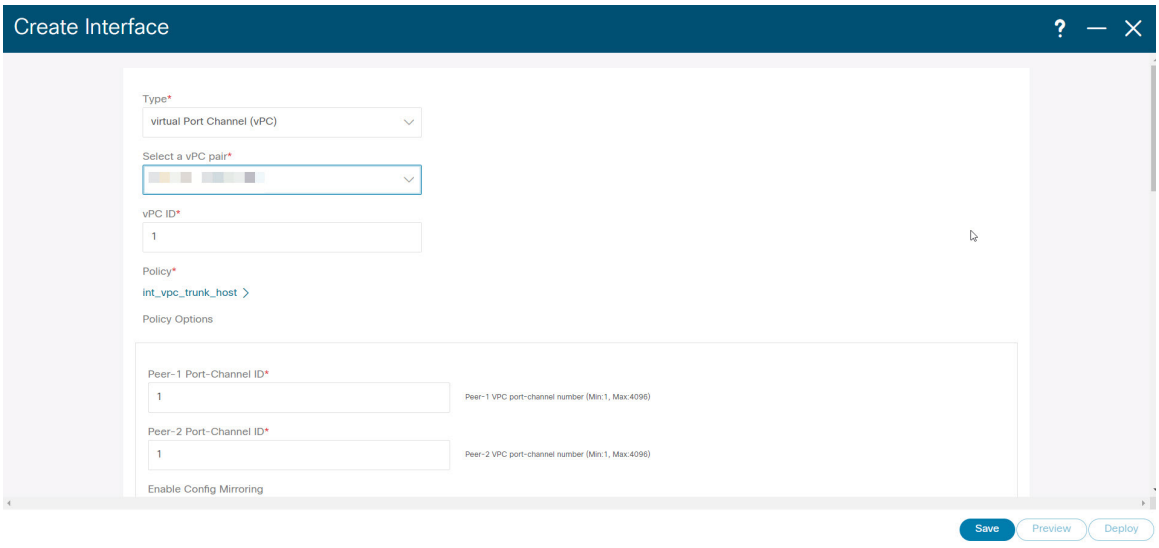


**Note** The setup of Data network interface on the Nexus Dashboard Fabric Controller is a prerequisite of any application that requires the in-band connectivity to the devices within fabric. This includes EPL and Network Insights Resources (NIR).

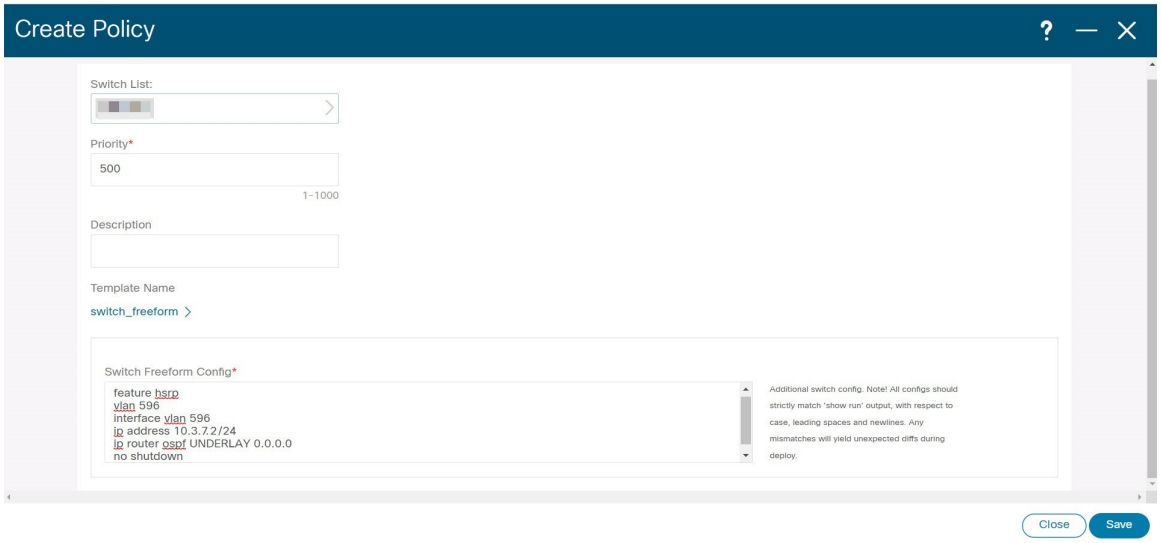
On the fabric side, for a standalone Nexus Dashboard Fabric Controller deployment, if the Nexus Dashboard data network port is directly connected to one of the front-end interfaces on a leaf, then that interface can be configured using the **epl_routed_intf** template. An example scenario of how this can be done when IS-IS or OSPF is employed as the IGP in the fabric, is depicted below:




However, for redundancy purposes, it is always advisable to have the server on which the Nexus Dashboard Fabric Controller is installed to be dual-homed or dual-attached. With the OVA Nexus Dashboard Fabric Controller deployment, the server can be connected to the switches via a port-channel. This provides link-level redundancy. To also have node-level redundancy on the network side, the server may be attached to a vPC pair of Leaf switches. In this scenario, the switches must be configured such that the HSRP VIP serves as the default gateway of the Data Network interface on the Nexus Dashboard Fabric Controller.



For the HSRP configuration on terry-leaf3, the **switch_freeform** policy may be employed as shown in the following image:




You can deploy a similar configuration on terry-leaf3 while using IP address 10.3.7.2/24 for SVI 596. This establishes an in-band connectivity from the Nexus Dashboard Fabric Controller to the fabrics over the Data Network interface with the default gateway set to 10.3.7.3.

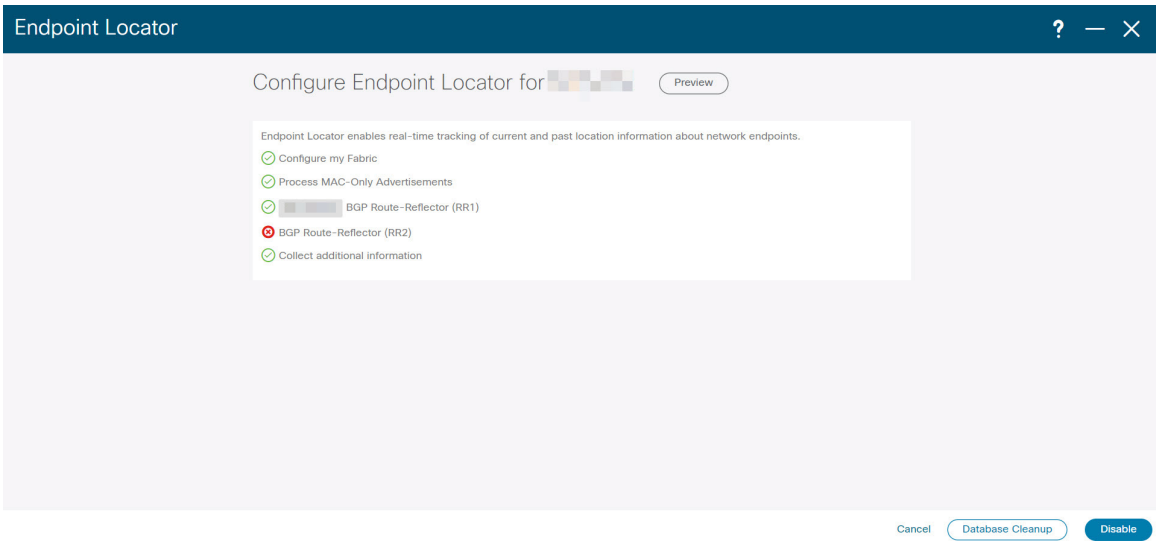
After you establish the in-band connectivity between the physical or virtual Nexus Dashboard Fabric Controller and the fabric, you can establish BGP peering.

During the EPL configuration, the route reflectors (RRs) are configured to accept Nexus Dashboard Fabric Controller as a BGP peer. During the same configuration, the Nexus Dashboard Fabric Controller is also configured by adding routes to the BGP loopback IP on the spines/RRs via the Data Network Interface gateway.

**Note** Ensure that you have enabled EPL feature for Cisco Nexus Dashboard Fabric Controller. Choose **Settings** > **Feature Management** > **Fabric Controller** choose check box **Endpoint Locator**. You can view the added EPL details on dashboard.

**Note** Cisco Nexus Dashboard Fabric Controller queries the BGP RR to glean information for establishment of the peering, like ASN, RR, IP, and so on.

To configure Endpoint Locator from the Cisco Nexus Dashboard Fabric Controller Web UI, On Fabric Overview page, choose **Actions** > **More** > **Configure Endpoint Locator**. Similarly, you can configure EPL on Topology page, right-click on required fabric, click **More** > **Configure Endpoint Locator**. The **Endpoint Locator** window appears.




You can enable EPL for one fabric at a time.

Select the switches on the fabric hosting the RRs from the drop-down list. Cisco Nexus Dashboard Fabric Controller will peer with the RRs.

By default, the **Configure My Fabric** option is selected. This option only configures EPL as a BGP neighbor of the switch and this option does not configure network reachability between EPL and the switch. This knob controls whether BGP configuration will be pushed to the selected spines/RRs as part of the enablement of the EPL feature. If the spine/RR needs to be configured manually with a custom policy for the EPL BGP neighborship, then this option should be unchecked. For external fabrics that are only monitored and not configured by Nexus Dashboard Fabric Controller, this option is greyed out as these fabrics are not configured by Nexus Dashboard Fabric Controller.

Select the **Process MAC-Only Advertisements** option to enable processing of MAC-Only advertisements while configuring the EPL feature.

**Note** If EPL is enabled on a fabric with or without selecting the **Process Mac-Only Advertisements** checkbox and you want to toggle this selection later, then you have to first disable EPL and then click **Database Clean-up** to delete endpoint data before re-enabling EPL with the desired **Process Mac-Only Advertisements** setting.

Select **Yes** under **Collect Additional Information** to enable collection of additional information such as PORT, VLAN, VRF etc. while enabling the EPL feature. To gather additional information, NX-API must be supported and enabled on the switches, ToRs, and leafs. If the **No** option is selected, this information will not be collected and reported by EPL.

**Note** For all fabrics except external fabrics, NX-API is enabled by default. For external fabrics, you have to enable NX-API in the external fabric settings by selecting the **Enable NX-API** checkbox in the **Advanced** tab of the External_Fabric_11_1 fabric template.

Click the **i** icon to view a template of the configuration that is pushed to the switches while enabling EPL. This configuration can be copied and pasted on spines or border gateway devices to enable EPL on external monitored fabrics.

Once the appropriate selections are made and various inputs have been reviewed, click **Submit** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled.

The Nexus Dashboard Data Service IP is used as BGP neighbor.

When the Endpoint Locator feature is enabled, there are a number of steps that occur in the background. Nexus Dashboard Fabric Controller contacts the selected RRs and determines the ASN. It also determines the interface IP that is bound to the BGP process. Also, appropriate BGP neighbor statements are added on the RRs or spines in case of eBGP underlay, to get them ready to accept the BGP connection that will be initiated from the Nexus Dashboard Fabric Controller. The external Nexus Dashboard Data Service IP address that is assigned to the EPL pod will be added as the BGP neighbor. Once EPL is successfully enabled, the user is automatically redirected to the EPL dashboard that depicts operational and exploratory insights into the endpoints that are present in the fabric.

For more information about the EPL dashboard, refer Monitoring Endpoint Locator, on page 180.

### Flushing the Endpoint Database

After you enable the Endpoint Locator feature, you can clean up or flush all the Endpoint information. This allows starting from a clean-slate with respect to ensuring no stale information about any endpoint is present in the database. After the database is clean, the BGP client re-populates all the endpoint information learnt from the BGP RR. You can flush the endpoint database even if you have not re-enabled the EPL feature on a fabric on which the EPL feature was previously disabled.

To flush all the Endpoint Locator information from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

#### Procedure

**Step 1** Choose **Endpoint Locator** > **Configure**, and click **Database Clean-Up**.

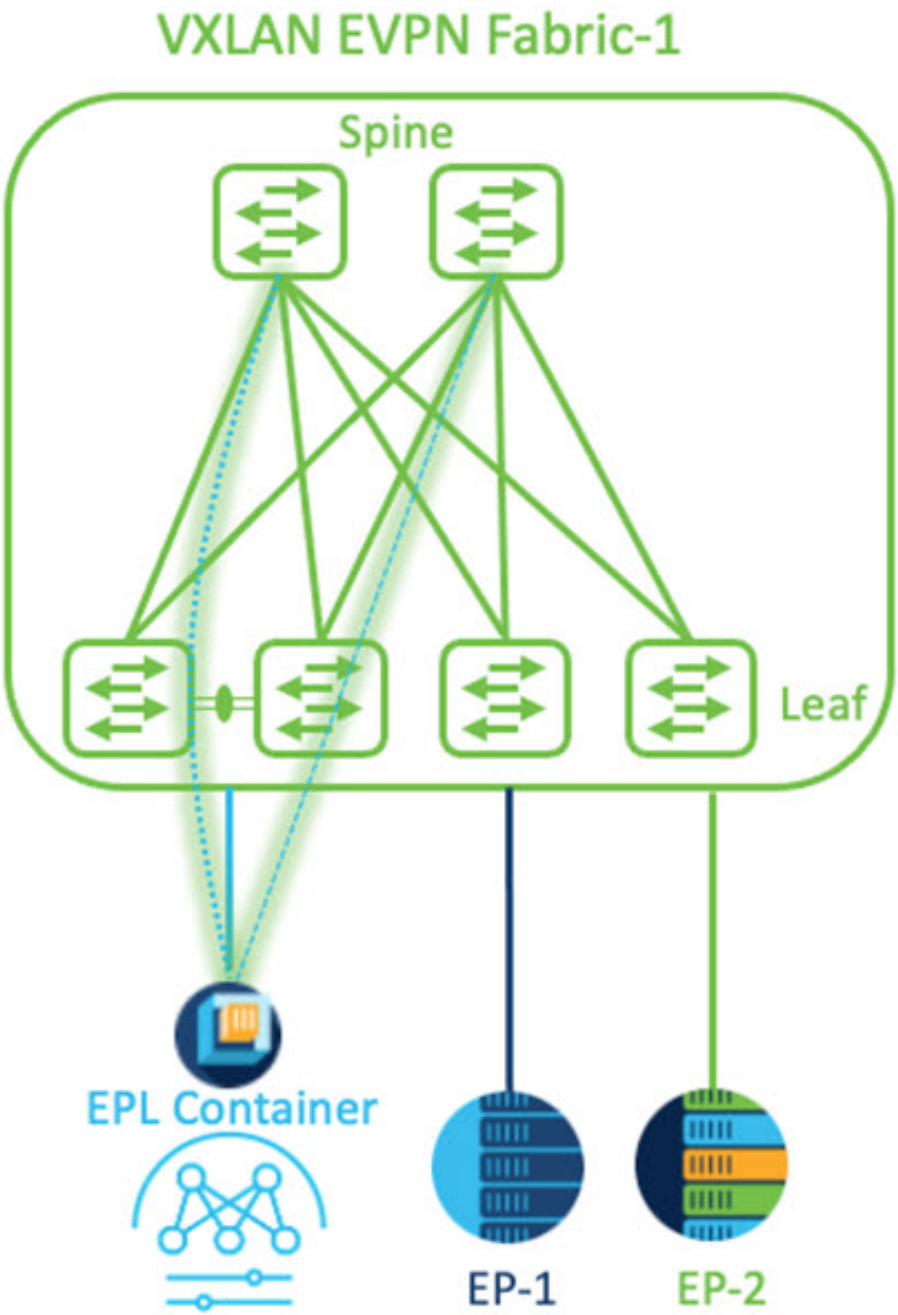
**Step 2** Click **Delete** to continue or **Cancel** to abort.

## Configuring Endpoint Locator for Single VXLAN EVPN Site

To configure endpoint locator for single VXLAN EVPN site, perform the following steps:

#### Before you begin

In the below figure, the NDFC service application is attached to the VPC pair of Leaf switches as it provides the link and node-level redundancy. The BGP instance running on EPL container establishes iBGP peering with the fabric spines. The iBGP peering is between Spine loopback addresses (loopback0) and EPL container persistent IP addresses. The loopback0 address of Spines is reachable via VXLAN Underlay, therefore, EPL container IP must have IP reachability towards the spines. We can configure an SVI on Leaf switches that can provide IP connectivity. The SVI will be a non-VXLAN enabled VLAN and will only participate in the underlay.




**Procedure**

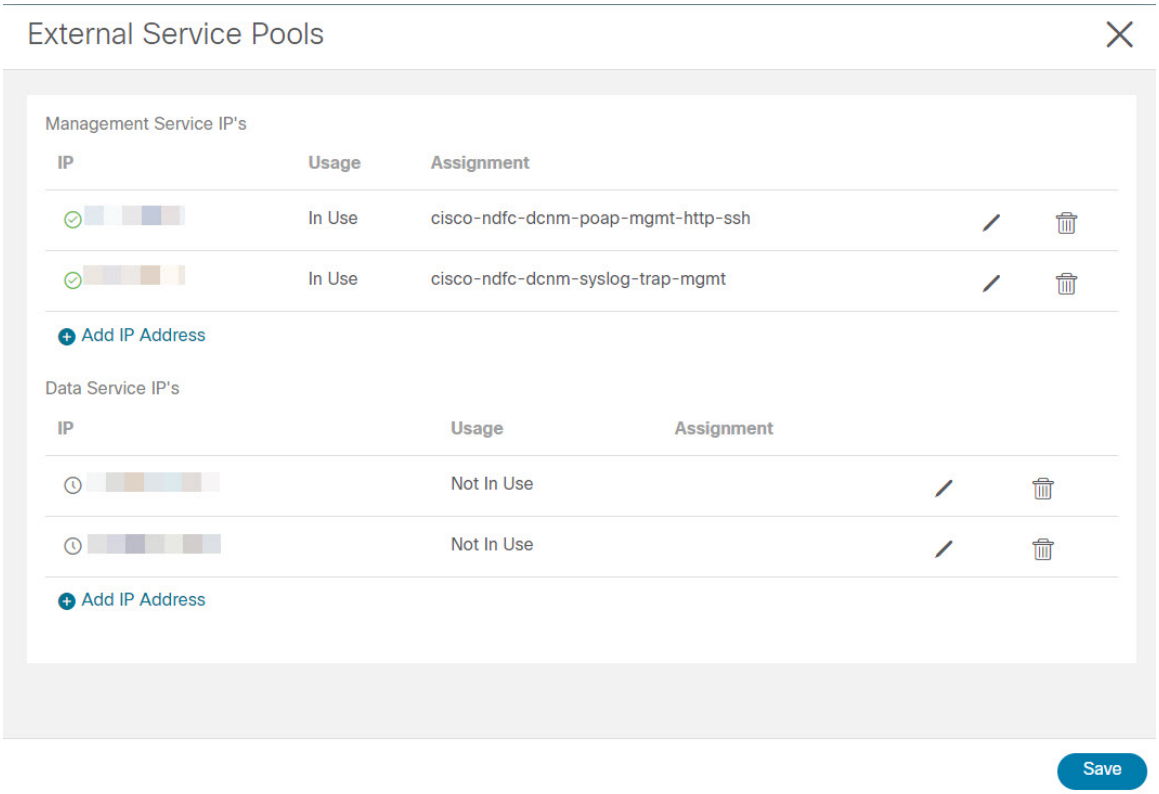
**Step 1** You must configure persistent IP addresses on Cisco Nexus Dashboard. On Nexus Dashboard, choose **Admin Console** > **Infrastructure** > **Cluster Configuration**.

**Step 2** On **General** tab, in **External Service Pools** card, click **Edit** icon.

The **External Service Pools** window appears.

**Step 3** Enter Persistent IP addresses in **Data Service IP's** and click **check** icon.

**Note**
The IP address must be associated with Nexus Dashboard Data Pool. A single persistent IP address is required to visualize and track EPs for a single site.




**Step 4** Configure SVI using FHRP for ND Data Interface and Underlay IP connectivity.

You can use **switch_freeform** policy on fabric Leaf 1.

To create a freeform policy, perform the following steps:

a) Choose **LAN**> **Fabrics**, double-click on required fabric.

The **Fabric Overview** page appears.

b) Click **Policy** tab, choose **Actions**> **Add Policy**.

The **Add Policy** window appears.

c) Choose appropriate Leaf1 switch from the **Switch List** drop-down list and click **Choose Template**.

d) On **Select Policy Template** window, choose **switch_freeform** template and click **Select**.

**Apply FHRP configurations and save the template.**

**Deploy the template configuration.**

In this example, SVI 100 with HSRP gateway created on fabric Leaf 1. Similarly, repeat the steps for fabric Leaf 2.

Below mentioned configuration example:

```
feature hsrp
vlan 100
name EPL-Inband
interface Vlan100
  no shutdown
  no ip redirects
  ip address 192.168.100.252/24
  no ipv6 redirects
  ip router ospf 100 area 0.0.0.0
  hsrp 100
    ip 192.168.100.254
```

**Step 5** Verify IP reachability between Nexus Dashboard Data Interface and fabric switches.

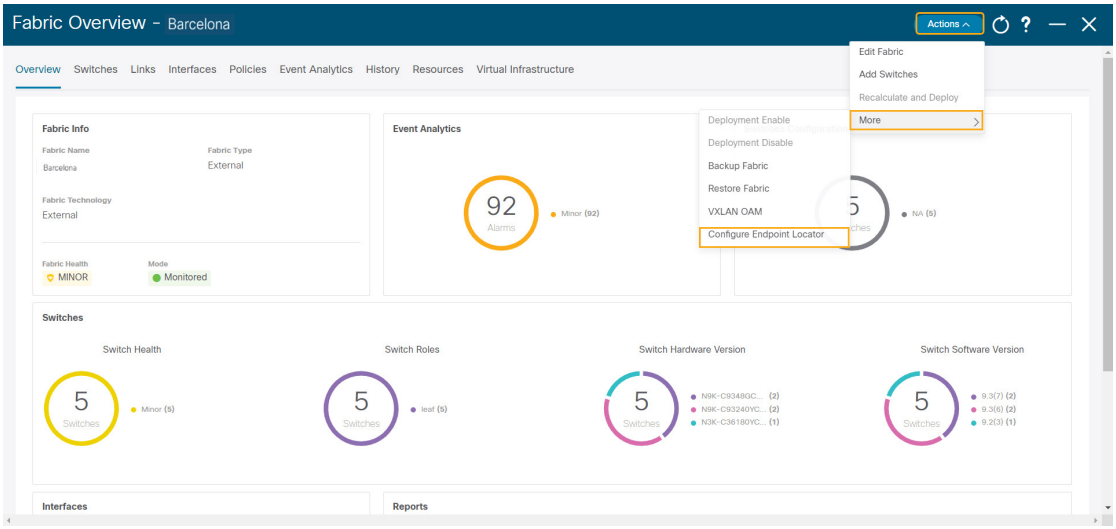
```
[rescue-user@ndfc-12-parth ~]$ ping 192.168.100.254 -c 2
PING 192.168.100.254 (192.168.100.254) 56(84) bytes of data.
64 bytes from 192.168.100.254: icmp_seq=1 ttl=255 time=1.95 ms
64 bytes from 192.168.100.254: icmp_seq=2 ttl=255 time=2.09 ms

--- 192.168.100.254 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 1.950/2.021/2.093/0.084 ms
[rescue-user@ndfc-12-parth ~]$
```

**Step 6** Enable EPL at fabric level.

a) To configure EPL, choose **LAN**> **Fabrics**> **Fabric Overview**.

b) On **Fabric Overview** window, choose **Actions**> **More**> **Configure EndPoint Locator**.




c) Choose the appropriate switches on the fabric hosting the Spine/Route Reflector RRs from the drop-down list.

Choose **Configure my Fabric** option for knob controls.

Whether BGP configuration will be pushed to the selected Spines/RRs as part of the enablement of the EPL feature. If the Spine/RR needs to be configured manually with a custom policy for the EPL BGP neighborship, then this option should be unchecked. For external fabrics that are only monitored and not configured on NDFC this option is grayed out. As these fabrics are not configured on NDFC.

Choose **Process MAC-Only Advertisements** option to enable processing of MAC-Only advertisements while configuring the EPL feature.

**Note**
If EPL is enabled on a fabric with or without choosing the **Process Mac-Only Advertisements** checkbox and if you want to toggle this selection later, then you must disable EPL and click **Database Clean-up** to delete endpoint data before re-enabling EPL with the desired **Process Mac-Only Advertisements** setting.

Choose **Yes** in **Collect Additional Information** to enable collection of additional information such as PORT, VLAN, and VRF while enabling the EPL feature. To access additional information, NX-API must be supported and enabled on the switches, ToRs, and leafs. If you choose **No** option, this information won't be collected and reported by EPL.

**Note**
For all fabrics except external fabrics, NX-API is enabled by default. For external fabrics, you must enable NX-API in the external fabric settings, choose **Enable NX-API** checkbox in the **Advanced** tab of the External_Fabric_11_1 fabric template.

Click on **Preview** icon to view a template of the configuration that is pushed to the switches enabling EPL. This configuration can be copied and pasted on spines or border gateway devices to enable EPL on external monitored fabrics.

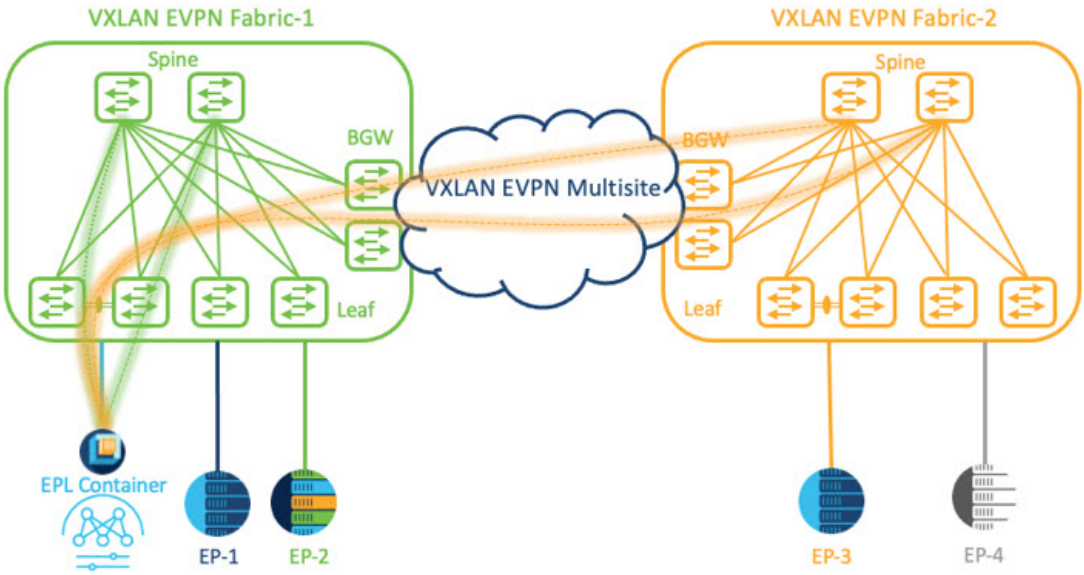
Once the appropriate selections are made and various inputs have been reviewed, click **Save Config** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message are displayed. Otherwise, EPL is successfully enabled. Once the EPL is enabled the Persistent IP will be in-use.

## Configuring Endpoint Locator for Multi-Fabric using VXLAN EVPN Multisite

To configure endpoint locator for mult-fabric VXLAN EVPN multisite, perform the following steps:

### Before you begin

The below figure enables EPL for Multi-Fabric using VXLAN EVPN Multisite. The BGP peering's are established between the Spines/RRs of each VXLAN EVPN Site and NDFC EPL Container. The Persistent IPs are required based on the number of VXLAN EVPN Sites. The NDFC application hosted on Cisco ND Cluster is located on Site 1. The routing information to reach the Spines/RRs deployed in the remote site must be exchanged across the Multisite. Once the BGP session is formed, local EPs of Fabric 2 can be visualized and tracked.




By default, Nexus Dashboard data Interface and Site 2 Spines/RRs loopback prefixes are not advertised across the BGWs. Therefore, prefixes must be exchanged using custom route maps and prefix lists across the sites. At the same time, route redistribution between OSPF and BGP is required as Spines/RRs loopback prefixes are part of OSPF protocol while BGWs peer with each other using BGP.

### Procedure

**Step 1** You must configure persistent IP addresses on Cisco Nexus Dashboard. On Nexus Dashboard, choose **Admin Console** > **Infrastructure** > **Cluster Configuration**.

**Step 2** On **General** tab, in **External Service Pools** card, click **Edit** icon.

The **External Service Pools** window appears.

**Step 3** Enter Persistent IP addresses in **Data Service IP's** and click **check** icon.

**Note**
Ensure that the IP addresses are associated with Nexus Dashboard Data Pool. Two persistent IP addresses are required to visualize and track EPs for a multisite with two member fabrics. One Persistent Data IP address is used as EPL container IP to establish BGP session with Site 1 fabric. A new Persistent IP address is configured that can be used to peer with Site 2 fabric.

**Step 4** Configure Route Redistribution for VXLAN EVPN Fabrics.

Route Redistribution for Fabric 1

The following switch_freeform policy can be used on Fabric 1 BGWs. To create a new **switch_freeform** policy, refer to the above examples.

Below the example of sample configuration

```
ip prefix-list site-2-rr seq 5 permit 20.2.0.1/32 >> Site 2 RR
ip prefix-list site-2-rr seq 6 permit 20.2.0.2/32 >> Site 2 RR
```

```
ip prefix-list epl-subnet seq 5 permit 192.168.100.0/24 >> EPL Subnet

route-map bgp-to-ospf permit 10
   match ip address prefix-list site-2-rr
route-map ospf-to-bgp permit 10
   match ip address prefix-list epl-subnet

router ospf 100
   redistribute bgp 100 route-map bgp-to-ospf

router bgp 100
   address-family ipv4 unicast
       redistribute ospf 100 route-map ospf-to-bgp
```

Route Redistribution for Fabric 2

The following switch_freeform policy can be used on Fabric 2 BGWs. To create a new **switch_freeform** policy, refer to the above examples.

Below the example of sample configuration

```
ip prefix-list site-2-rr seq 5 permit 20.2.0.1/32 >> Site 2 RR
ip prefix-list site-2-rr seq 6 permit 20.2.0.2/32 >> Site 2 RR
ip prefix-list epl-subnet seq 5 permit 192.168.100.0/24 >> EPL Subnet

route-map bgp-to-ospf permit 10
  match ip address prefix-list epl-subnet
route-map ospf-to-bgp permit 10
  match ip address prefix-list site-2-rr

router ospf 200
  redistribute bgp 200 route-map bgp-to-ospf

router bgp 200
    address-family ipv4 unicast
        redistribute ospf 200 route-map ospf-to-bgp
```

**Step 5** To configure EPL, choose **LAN**> **Fabrics**> **Fabric Overview**.

**Step 6** On **Fabric Overview** window, choose **Actions**> **More**> **Configure EndPoint Locator**.

**Step 7** Choose the appropriate switches on the fabric hosting the Spine/Route Reflector RRs from the drop-down list.

Once the appropriate selections are made and various inputs have been reviewed, click **Save Config** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled. Once the EPL is enabled the Persistent IP will be in-use.

You can view EPL enabled for fabric-1 and fabric-2 successfully. To view and track EPs, Refer the Monitoring Endpoint Locator section.

## Configuring Endpoint Locator for vPC Fabric Peering Switches

Networks Administrator can create vPC between a pair of switches using a Physical Peer Link or Virtual Peer link. vPC Fabric Peering provides an enhanced dual-homing access solution without the overhead of wasting physical ports for vPC Peer Link. For Virtual Peer link, EPL can still be connected to vPC pair of Leaf switches for the link and node-level redundancy. However, VXLAN VLAN (Anycast Gateway) as the First hop for EPL will be used. The loopback0 address of Spines/RRs is reachable only via VXLAN Underlay, while VXLAN VLAN will be part of a Tenant VRF. Therefore, to establish IP communication, route-leaking is configured between Tenant VRF and Default VRF. For more information, refer to vPC Fabric Peering section.

To configure endpoint locator for vPC Fabric Peering switches perform the following steps:

### Procedure

**Step 1** You must configure persistent IP addresses on Cisco Nexus Dashboard. On Nexus Dashboard, choose **Admin Console** > **Infrastructure** > **Cluster Configuration**.

**Step 2** On **General** tab, in **External Service Pools** card, click **Edit** icon.

The **External Service Pools** window appears.

**Step 3** Enter Persistent IP addresses in **Data Service IP's** and click **check** icon.

**Step 4** Create a Tenant VRF and Anycast Gateway on the vPC fabric peering switches.

add two images

**Step 5** Configure Route-leaking between Tenant VRF and Default VRF.

**Advertise from Tenant VRF to Default VRF.**

The following switch_freeform policy can be used on fabric Leaf where ND is connected.

```
ip prefix-list vrf-to-default seq 5 permit 192.168.100.0/24 >> EPL subnet
route-map vrf-to-default permit 10
  match ip address prefix-list vrf-to-default
vrf context epl_inband
  address-family ipv4 unicast
    export vrf default map vrf-to-default allow-vpn
router ospf UNDERLAY
   redistribute bgp 200 route-map vrf-to-default
```

**Advertise from Default VRF to Tenant VRF.**

The following switch_freeform policy can be used on fabric Leaf where ND is connected.

```
ip prefix-list default-to-vrf seq 5 permit 20.2.0.3/32 >> Spine loopback IP
ip prefix-list default-to-vrf seq 6 permit 20.2.0.4/32 >> Spine loopback IP
route-map default-to-vrf permit 10
   match ip address prefix-list default-to-vrf
vrf context epl_inband
   address-family ipv4 unicast
       import vrf default map default-to-vrf
        router bgp 200
   address-family ipv4 unicast
       redistribute ospf UNDERLAY route-map default-to-vrf
```

**Step 6** Enable EPL at fabric level.

a) To configure EPL, choose **LAN**> **Fabrics**> **Fabric Overview**.

b) On **Fabric Overview** window, choose **Actions**> **More**> **Configure EndPoint Locator**.

c) Choose the appropriate switches on the fabric hosting the Spine/Route Reflector RRs from the drop-down list.

Once the appropriate selections are made and various inputs have been reviewed, click **Save Config** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled. Once the EPL is enabled the Persistent IP will be in-use.

# Configuring Endpoint Locator for External Fabrics

In addition to Easy fabrics, Nexus Dashboard Fabric Controller allows you to enable EPL for VXLAN EVPN fabrics comprising of switches that are imported into the external fabric. The external fabric can be in managed mode or monitored mode, based on the selection of **Fabric Monitor Mode** flag in the **External Fabric** Settings. For external fabrics that are only monitored and not configured by Nexus Dashboard Fabric Controller, this flag is disabled. Therefore, you must configure BGP sessions on the Spine(s) via OOB or using the CLI. To check the sample template, click  icon to view the configurations required while enabling EPL.

In case the **Fabric Monitor Mode** checkbox in the External Fabric settings is unchecked, then EPL can still configure the spines/RRs with the default **Configure my fabric** option. However, disabling EPL would wipe out the router bgp config block on the spines/RRs. To prevent this, the BGP policies must be manually created and pushed onto the selected spines/RRs.

# Configuring Endpoint Locator for eBGP EVPN Fabrics

You can enable EPL for VXLAN EVPN fabrics, where eBGP is employed as the underlay routing protocol. Note that with an eBGP EVPN fabric deployment, there is no traditional RR similar to iBGP. The reachability of the in-band subnet must be advertised to the spines that behave as Route Servers. To configure EPL for eBGP EVPN fabrics from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **LAN > Fabrics**.

Select the fabric to configure eBGP on or create eBGP fabric with the **Easy_Fabric_eBGP** template.

**Step 2** Use the **leaf_bgp_asn** policy to configure unique ASNs on all leaves.

**Step 3** Add the **ebgp_overlay_leaf_all_neighbor** policy to each leaf.

Fill **Spine IP List** with the spines' BGP interface IP addresses, typically the loopback0 IP addresses.

Fill **BGP Update-Source Interface** with the leaf's BGP interface, typically loopback0.

**Step 4** Add the **ebgp_overlay_spine_all_neighbor** policy to each spine.

Fill **Leaf IP List** with the leaves' BGP interface IPs, typically the loopback0 IPs.

Fill **Leaf BGP ASN** with the leaves' ASNs in the same order as in **Leaf IP List**.

Fill **BGP Update-Source Interface** with the spine's BGP interface, typically loopback0.

After the in-band connectivity is established, the enablement of the EPL feature remains identical to what is listed so far. EPL becomes a iBGP neighbor to the Route Servers running on the spines.

# Monitoring Endpoint Locator

Information about the Endpoint Locator is displayed on a single landing page or dashboard. The dashboard displays an almost real-time view of data (refreshed every 30 seconds) pertaining to all the active endpoints

on a single pane. The data that is displayed on this dashboard depends on the scope selected by you from the **SCOPE** drop-down list. The Nexus Dashboard Fabric Controller scope hierarchy starts with the fabrics. Fabrics can be grouped into a Multi-Site Domain (MSD). A group of MSDs constitute a Data Center. The data that is displayed on the Endpoint Locator dashboard is aggregated based on the selected scope. From this dashboard, you can access Endpoint History, Endpoint Search, and Endpoint Life.



**Note** This is a preview feature in Nexus Dashboard Fabric Controller, Release 12.0.1a. We recommend that you use this feature marked as BETA in your lab setup only. Do not use these features in your production deployment.

# Disabling Endpoint Locator

To disable endpoint locator from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **Endpoint Locator > Configure**.

The **Endpoint Locator** window appears. Select the required fabric from the **SCOPE** dropdown list. The fabric configuration details are then displayed for the selected fabric.

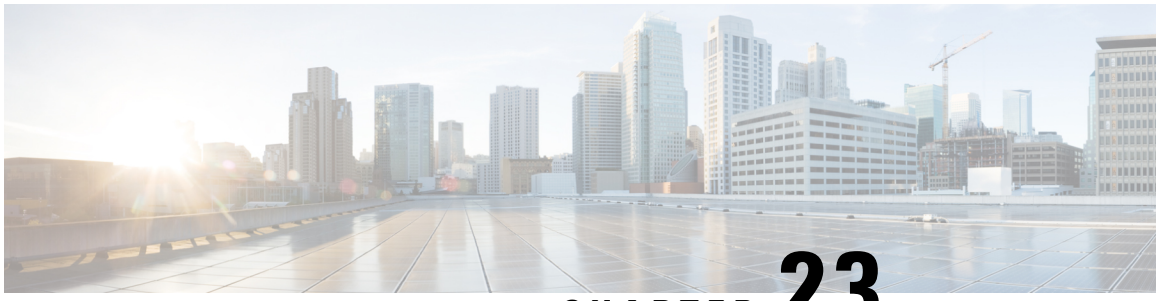
**Step 2** Click **Disable**.



PART VII

# Easy Provisioning of VXLAN BGP EVPN Fabrics


- Managing a Greenfield VXLAN BGP EVPN Fabric, on page 525
- Managing a Brownfield VXLAN BGP EVPN Fabric , on page 551
- Configuring a VXLANv6 Fabric, on page 585
- Multi-Site Domain for VXLAN BGP EVPN Fabrics , on page 589
- Configuring ToR Switches and Deploying Networks in External Fabrics, on page 607



CHAPTER 23

# Managing a Greenfield VXLAN BGP EVPN Fabric

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


- Provisioning VXLAN EVPN Fabric with IGP Underlay, on page 525
- Provisioning VXLAN EVPN Fabric with eBGP Underlay, on page 539


# Provisioning VXLAN EVPN Fabric with IGP Underlay

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

For creating and deploying VXLAN EVPN fabrics, see VXLAN EVPN Fabrics Provisioning, on page 45.

## Creating VXLAN EVPN Fabric with IPv4 Underlay

To create a new VXLAN EVPN fabric, refer to Creating a VXLAN EVPN Fabric Using the Easy_Fabric Template, on page 49.

## Creating VXLAN EVPN Fabric with IPv6 Underlay

This procedure shows how to create a VXLAN EVPN fabric with IPv6 underlay. Note that only the fields for creating a VXLAN fabric with IPv6 underlay are documented. For information about the remaining fields, see Creating a VXLAN EVPN Fabric Using the Easy_Fabric Template, on page 49.

### Procedure

**Step 1** Choose **LAN** > **Fabrics**.

**Step 2** From the **Actions** drop-down list, choose **Create Fabric**.

The **Create Fabric** window appears.

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

**Fabric Template** – From the drop-down list, choose **Data Center VXLAN EVPN**.

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

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

**Enable IPv6 Underlay** – Check the **Enable IPv6 Underlay** check box .

**Enable IPv6 Link-Local Address** – Check the **Enable IPv6 Link-Local Address** check box to use the link local addresses in the fabric between leaf-spine and spine-border interfaces. If you check this check box, the **Underlay Subnet IPv6 Mask** field is not editable. By default, the **Enable IPv6 Link-Local Address** field is enabled.

IPv6 underlay supports **p2p** networks only. Therefore, the **Fabric Interface Numbering** drop-down list is disabled.

**Underlay Subnet IPv6 Mask** – Specify the subnet mask for the fabric interface IPv6 addresses.

**Underlay Routing Protocol** – Specify the IGP used in the fabric, that is, OSPF or IS-IS for VXLANv6.

**Step 4** All the fields under the **Replication** tab are disabled.

IPv6 underlay supports ingress replication mode only.

**Step 5** Click the **VPC** tab.

**vPC Peer Keep Alive option** – Choose **management** or **loopback**. To use IP addresses assigned to the management port and the management VRF, choose management. To use IP addresses assigned to loopback interfaces and a non-management VRF, choose underlay routing loopback with IPv6 address for PKA. Both the options are supported for IPv6 underlay.

**Step 6** Click the **Protocols** tab.

**Underlay Anycast Loopback Id** – Specify the underlay anycast loopback ID for IPv6 underlay. You cannot configure IPv6 address as secondary, an additional loopback interface is allocated on each vPC device. Its IPv6 address is used as the VIP.

**Step 7** Click the **Resources** tab.

**Manual Underlay IP Address Allocation**: Check the check box to manually allocate underlay IP addresses. The dynamic underlay IP addresses fields are disabled.

**Underlay Routing Loopback IPv6 Range**: Specify loopback IPv6 addresses for protocol peering.

**Underlay VTEP Loopback IPv6 Range**: Specify loopback IPv6 addresses for VTEPs.

**Underlay Subnet IPv6 Range**: Specify the IPv6 address range that is used for assigning IP addresses for numbered and peer link SVIs. To edit this field, uncheck **Enable IPv6 Link-Local Address** check box under the **General Parameters** tab.

**BGP Router ID Range for IPv6 Underlay**: Specify the address range to assign BGP Router IDs. The IPv4 addressing is used for router with BGP and underlay routing protocols.

**Step 8** Click the **Bootstrap** tab.

**Enable Bootstrap**: Check the **Enable Bootstrap** check box. If this check box is not chosen, none of the other fields on this tab are editable.

**Enable Local DHCP Server**: Check the check box to initiate automatic assignment of IP addresses assignment through the local DHCP server. The **DHCP Scope Start Address** and **DHCP Scope End Address** fields are editable only after you check this check box.

**DHCP Version**: Choose DHCPv4 from the drop-down list.

**Step 9** Click **Save** to complete the creation of the fabric.

### What to do next

Adding Switches to a Fabric

# Adding Switches

Switch can be added to a single fabric at any point in time. To add switches to a fabric and discover existing or new switches, refer to Adding Switches to a Fabric, on page 279.

# Assigning Switch Roles

To assign roles to switches on Nexus Dashboard Fabric Controller refer to Assigning Switch Roles, on page 297.

# Creating vPC Setup

(Optional) Create a vPC setup for a pair of switches in the fabric. Ensure that the switches have the same roles and are connected to each other. For instructions, refer to vPC Fabric Peering, on page 73.

# Overlay Mode

You can create a VRF or network in CLI or config-profile mode at the fabric level. The overlay mode of member fabrics of an MSD fabric is set individually at the member-fabric level. Overlay mode can only be changed before deploying overlay configurations to the switches. After the overlay configuration is deployed, you cannot change the mode unless all the VRF and network attachments are removed.

**Note** If you upgrade from Cisco DCNM Release 11.5(x), the existing config-profile mode functions the same.

If the switch has config-profile based overlays, you can import it in the **config-profile** overlay mode only. If you import it in the **cli** overlay mode, an error appears during brownfield import.

For brownfield import, if overlay is deployed as **config-profile** mode, it can be imported in **config-profile** mode only. However, if overlay is deployed as **cli**, it can be imported in either **config-profile** or **cli** modes.

To choose the overlay mode of VRFs or networks in a fabric, perform the following steps:

1. Navigate to the **Edit Fabric** window.
2. Go to the **Advanced** tab.
3. From the **Overlay Mode** drop-down list, choose **config-profile** or **cli**.

The default mode is **config-profile**.

## Creating VRF

**UI Navigation**

The following options are applicable only for switch fabrics, Easy fabrics, and MSD fabrics.

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **VRFs** > **VRFs**.
- Choose **LAN** > **Fabrics**. Double-click on the fabric to open **Fabric Overview** > **VRFs** > **VRFs**.

To create VRF from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

**Procedure**

**Step 1** On the **VRFs** tab, click **Actions** > **Create**.

The **Create VRF** window appears.

**Step 2** On **Create VRF**, enter the required details in the mandatory fields. The available fields vary based on the fabric type.

The fields in this window are:

**VRF Name** – Specifies a VRF name automatically or allows you to enter a name. The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).

For MSD Fabrics, the values for VRF or Network is same for the fabric.

**VRF ID** – Specifies the ID for the VRF or allows you to enter an ID for the VRF.

**VLAN ID** – Specifies the corresponding tenant VLAN ID for the network or allows you to enter an ID for the VLAN. If you want to propose a new VLAN for the network, click **Propose VLAN**.

**VRF Template** – A default universal template is auto-populated. This is applicable for leaf switches only.

**VRF Extension Template** – A default universal extension template is auto-populated. 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.

**Step 3** The fields on the **General** tab are:

**VRF VLAN Name** – Enter the VLAN name for the VRF.

**VRF Interface Description** – Enter a description for the VRF interface.

**VRF Description** – Enter a description for the VRF.

**Step 4** Click the **Advanced** tab to optionally specify the advanced profile settings. The fields on this tab are auto-populated. The fields on the **Advanced** tab are:

**VRF Interface MTU** – Specifies VRF interface MTU.

**Loopback 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 also.

**Redistribute Direct Route Map** – Specifies the redistribute direct route map name.

**Max BGP Paths** – Specifies the maximum number of BGP paths. The valid value is between 1 and 64.

**Max iBGP Paths** – Specifies the maximum number of iBGP paths. The valid value is between 1 and 64.

**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.

**TRM Enable** – Check the check box to enable TRM.

If you enable TRM, and provide the RP address, you must enter the underlay multicast address in the **Underlay Mcast Address**.

**Is RP External** – Check this check box if the RP is external to the fabric. If this check box is not checked, 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 TRM BGW MSite** – Check the check box to enable TRM on Border Gateway Multisite.

**Advertise Host Routes** – Check this check box to control advertisement of /32 and /128 routes to Edge routers.

**Advertise Default Route** – Check this check box to control advertisement of default route 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** check box) for the associated VRF. This will result in /32 routes for hosts 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 one fabric only then the default route is sufficient for inter-subnet communication.

**Config Static 0/0 Route** – Check this check box to control configuration of static default route.

**BGP Neighbor Password** – Specifies the VRF Lite BGP neighbor password.

**BGP Password Key Encryption Type** – From the drop-down list, select the encryption type.

**Enable Netflow** – Allows you to enable netflow monitoring on the VRF-Lite sub-interface. Note that this is supported only if netflow is enabled on the fabric.

**Netflow Monitor** – Specifies the monitor for the VRF-lite netflow configuration.

To enable netflow on a VRF-Lite sub-interface, you must enable netflow at VRF level and VRF extension level. Check the **Enable_IFC_Netflow** check box in the VRF attachment while you edit an extension to enable netflow monitoring.

For more information, refer to Netflow Support, on page 149.

**Step 5** Click **Create** to create the VRF or click **Cancel** to discard the VRF.

A message appears indicating that the VRF is created.

The new VRF appears on the **VRFs** horizontal tab. The status is **NA** as the VRF is created but not yet deployed. Now that the VRF is created, you can create and deploy networks on the devices in the fabric.

## VRF Attachments

**UI Navigation**

The following options are applicable only for switch fabrics, Easy fabrics, and MSD fabrics.

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **VRFs** > **VRF Attachments**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **VRFs** > **VRF Attachments**.

Use this window to attach or detach attachments to or from a VRF, respectively. You can also import or export the attachments for a VRF.

***Table 49: VRF Attachments Table Fields and Description***

| Field | Description |
|---|---|
| VRF Name | Specifies the name of the VRF. |
| VRF ID | Specifies the ID of the VRF. |
| VLAN ID | Specifies the VLAN ID. |
| Switch | Specifies the name of the switch. |
| Status | Specifies the status of VRF attachments, for example, pending, NA, deployed, out-of-sync, and so on. |
| Attachment | Specifies whether the VRF attachment is attached or detached. |
| Switch Role | Specifies the switch role. For example, for the fabric created using the Easy_Fabric_IOS_XE fabric template, the switch role is specified as either leaf, spine, or border. |
| Fabric Name | Specifies the name of the fabric to which the VRF is attached or detached. |
| Loopback ID | Specifies the loopback ID. |
| Loopback IPV4 Address | Specifies the loopback IPv4 address. |
| Loopback IPV6 Address | Specifies the loopback IPv6 address.<br>**Note**<br>The IPv6 address is not supported for underlay. |

Click the table header to sort the entries in alphabetical order of that parameter.

The following table describes the action items, in the **Actions** drop-down list, that appears on the **VRF Attachments** horizontal tab of the **VRFs** tab in the **Fabric Overview** window.

***Table 50: VRF Attachments Actions and Description***

| Action Item | Description |
|---|---|
| History | Allows you to view the deployment and policy change history of the selected VRF.<br><br>You can view the deployment history details of a VRF attachment such as hostname, VRF name, commands, status, status description, user, and completed time on the **Deployment History** tab.<br><br>You can view the policy change history details such as policy ID, template, description, PTI operation, generated configuration, entity name and type, created date, serial number, user, and source of the policy on the **Policy Change History** tab.<br><br>To view the history of a VRF attachment, check the check box next to the VRF name and select **History**. The **History** window appears. Click the **Deployment History** or **Policy Change History** tabs as required. You can also click the **Detailed History** link in the **Commands** column of the **Deployment History** tab to view the command execution details (comprising configuration, status, and CLI response) for the host. |
| Edit | Allows you to view or edit the VRF attachment parameters such as interfaces that you want to attach to the selected VRF.<br><br>To edit the VRF attachment information, check the check box next to the VRF name that you want to edit. Select **Edit**. In the **Edit VRF Attachment** window, edit the required values, attach or detach the VRF attachment. Click the **Edit** link to edit the CLI freeform config for the switch, and click **Save** to apply the changes or click **Cancel** to discard the changes. The edited VRF attachment is shown in the table on the **VRF Attachments** horizontal tab of the **VRFs** tab in the **Fabric Overview** window. |

| Action Item | Description |
|---|---|
| Preview | Allows you to preview the configuration of the VRF attachments for the selected VRF.<br>**Note**<br>This action is not allowed for attachments that are in deployed or NA status.<br><br>To preview the VRF, check the check box next to the VRF name and choose **Preview** from **Actions** drop-down list. The **Preview Configuration** window for the fabric appears.<br><br>You can preview the VRF attachment details such as the VRF name, fabric name, switch name, serial number, IP address, and role, VRF status, pending configuration, and progress of the configuration. Click the lines link in the **Pending Config** column to view the lines for which the configuration is pending. Click **Close**. |
| Deploy | Allows you to deploy the pending configuration of the VRF attachments, for example, interfaces, for the selected VRF.<br>**Note**<br>This action is not allowed for attachments that are in deployed or NA status.<br><br>To deploy a VRF, check the check box next to the VRF name and choose **Deploy** from **Actions** drop-down list. The **Deploy Configuration** window for the fabric appears.<br><br>You can view the details such as the VRF name, fabric name, switch name, serial number, IP address, and role, VRF status, pending configuration, and progress of the configuration. Click the lines link in the **Pending Config** column to view the lines for which the configuration is pending. Click the **Deploy** button. The status and progress of the deployment is displayed in the **VRF Status** and **Progress** columns. After the deployment is completed successfully, close the window. |
| Import | Allows you to import information about VRF attachments for the selected fabric.<br><br>To import the VRF attachments information, choose **Import**. Browse the directory and select the `.csv` file that contains the VRF attachments information. Click **Open** and then click **OK**. The VRF information is imported and displayed in the **VRF Attachments** horizontal tab on the **VRFs** tab in the **Fabric Overview** window. |

| Action Item | Description |
|---|---|
| Export | Allows you to export the information about VRF attachments to a `.csv` file. The exported file contains information pertaining to each VRF, including the fabric it belongs to, whether the LAN is attached, the associated VLAN, serial number, interfaces, and freeform configuration details that you saved for VRF attachments.<br>To export VRF attachments information, choose the **Export** action. Select a location on your local system directory to store the VRF information and click **Save**. The VRF information file is exported to your local directory. The file name is appended with the date and time at which the file was exported. |
| Quick Attach | Allows you to immediately attach an attachment to the selected VRF. You can select multiple entries and attach them to a VRF at the same instance.<br>To quickly attach any attachment to a VRF, choose **Quick Attach** from **Actions** drop-down list. A message appears to inform that the attach action was successful. |
| Quick Detach | Allows you to detach the selected VRF immediately from an attachment, for example, a fabric. You can select multiple entries and detach them from an attachment at the same instance.<br>To attach any attachment to a VRF quickly, choose **Quick Detach** from **Actions** drop-down list. A message appears to inform that the detach action was successful. |

## Creating Network for Standalone Fabrics

To create a network from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

### Before you begin

Before creating networks, ensure that you have created a VRF for the fabric. However, if you have chosen Layer 2 on the **Create Network** window, then you do not require a VRF. For more information, see VRFs, on page 203.

### Procedure

**Step 1** On the **Networks** tab, click **Actions** > **Create**.

The **Create Network** window appears.

**Step 2** On **Create Network**, enter the required details in the mandatory fields. The available fields vary based on the fabric type.

**Note**

If the fields for the **Network ID** field below and the **VRF ID** field (after clicking **Create VRF**) are not automatically populated, one possible reason is that the VNI ranges might be exhausted. In this situation, you can extend the range for VNI accordingly in **Fabric Settings**.

The fields in this window are:

**Network ID** and **Network Name** – Specifies the Layer 2 VNI and the 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.

**Layer 2 Only** – Specifies whether the network is Layer 2 only.

**VRF Name** – Allows you to select the Virtual Routing and Forwarding (VRF) from the drop-down list.

If you want to create a new VRF, click **Create VRF**. The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).

**VLAN ID** – Specifies the corresponding tenant VLAN ID for the network. If you want to propose a new VLAN for the network, click **Propose VLAN**.

**Network Template** – A default universal template is auto-populated. This is only applicable for leaf switches.

**Network Extension Template** – A default universal extension template is auto-populated. 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.

**Generate Multicast IP** – Click to generate a new multicast group address and override the default value.

**Step 3** The fields on the **General Parameters** tab are:

**Note**
If the network is a non Layer 2 network, then it is mandatory to provide the gateway IP address.

**IPv4 Gateway/NetMask**: Specifies the IPv4 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. The anycast gateway IP address is the same for MyNetwork_30000 on all switches of the fabric that have the presence of the network.

**Note**
If the same IP address is configured in the IPv4 Gateway and IPv4 Secondary GW1 or GW2 fields of the network template, Nexus Dashboard Fabric Controller 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 List** – Specifies the IPv6 address with subnet.

**Vlan Name** – Enter the VLAN name.

**Interface Description** – Specifies the description for the interface. This interface is a switch virtual interface (SVI).

**MTU for L3 interface** – Enter the MTU for Layer 3 interfaces range 68 - 9216.

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

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

**IPv4 Secondary GW3** – Enter the gateway IP address for the additional subnet.

**IPv4 Secondary GW4** – Enter the gateway IP address for the additional subnet.

**Step 4** Click the **Advanced** tab to optionally specify the advanced profile settings. The fields on the **Advanced** tab are:

**ARP Suppression** – Select the check box to enable the ARP Suppression function.

**Ingress Replication** – The check box 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 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.

Starting from Cisco NDFC Release 12.1.2e, a maximum of 16 DHCP relay servers for overlay networks are supported. Perform the following steps to include the DHCP relay server information:

**a.** a. On the **DHCP Relay Server Information** field, click **Actions** > **Add**.

The **ADD Item** window appears.

**b.** Enter the **Server IP V4 Address** and **Server VRF** details and click **Save**.

**c.** Repeat the above steps to add the required number of DHCP relay server information.

**DHCPv4 Server 1** – Enter the DHCP relay IP address of the first DHCP server.

**DHCPv4 Server VRF** – Enter the DHCP server VRF ID.

**DHCPv4 Server 2** – Enter the DHCP relay IP address of the next DHCP server.

**DHCPv4 Server2 VRF** – Enter the DHCP server VRF ID.

**DHCPv4 Server 3** – Enter the DHCP relay IP address of the next DHCP server.

**DHCPv4 Server3 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** – Check the check box to enable TRM.

For more information, see Overview of Tenant Routed Multicast.

**L2 VNI Route-Target Both Enable** – Check the check box to enable automatic importing and exporting of route targets for all L2 virtual networks.

**Enable Netflow** – Enables netflow monitoring on the network. This is supported only if netflow is already enabled on fabric.

**Interface Vlan Netflow Monitor** – Specifies the netflow monitor specified for Layer 3 record for the VLAN interface. This is applicable only if **Is Layer 2 Record** is not enabled in the **Netflow Record** for the fabric.

**Vlan Netflow Monitor** – Specifies the monitor name defined in the fabric setting for Layer 3 **Netflow Record**.

**Enable L3 Gateway on Border** – Check the check box to enable a Layer 3 gateway on the border switches.

**Step 5** Click **Create**.

A message appears indicating that the network is created.

The new network appears on the **Networks** page that comes up.

The Status is **NA** since the network is created but not yet deployed on the switches. Now that the network is created, you can create more networks if necessary and deploy the networks on the devices in the fabric.

## Network Attachments

**UI Navigation**

The following options are applicable only for switch fabrics, Easy fabrics, and MSD fabrics:

- Choose **LAN** > **Fabrics**. Click on the fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Networks** > **Network Attachments**.
- Choose **LAN** > **Fabrics**. Double-click on the fabric to open **Fabric Overview** > **Networks** > **Network Attachments**.

Use this window to attach fabrics and interfaces to a network.

***Table 51: Network Attachments Table Fields and Description***

| Field | Description |
|---|---|
| Network Name | Specifies the name of the network. |
| Network ID | Specifies the Layer 2 VNI of the network. |
| VLAN ID | Specifies the VLAN ID. |
| Switch | Specifies the name of the switch. |
| Ports | Specifies the ports for the interfaces. |
| Status | Specifies the status of the network attachments, for example, pending, NA, and so on. |
| Attachment | Specifies whether the network attachment is attached or detached. |
| Switch Role | Specifies the switch role. For example, for the fabric created using the Easy_Fabric_IOS_XE fabric template, the switch role is specified as either leaf, spine, or border. |
| Fabric Name | Specifies the name of the fabric to which the network is attached or detached. |

The following table describes the action items, in the **Actions** drop-down list, that appears in the **Network Attachments** horizontal tab on the **Networks** tab in the **Fabric Overview** window.

*Table 52: Network Attachments Actions and Description*

| Action Item | Description |
|---|---|
| History | Allows you to view the deployment and policy change history of the selected network.<br><br>You can view the deployment history details of a network attachment such as hostname, network name, VRF name, commands, status, status description, user and completed time on the **Deployment History** tab.<br><br>You can view the policy change history details such as policy ID, template, description, PTI operation, generated configuration, entity name and type, created date, serial number, user, and source of the policy on the **Policy Change History** tab.<br><br>To view the history of a network attachment, select the check box next to the network name and choose the **History** action. The **History** window appears. Click the **Deployment History** or **Policy Change History** tabs as required. Click the **Detailed History** link in the **Commands** column of the **Deployment History** tab to view the command execution details (comprising configuration, status, and CLI response) for the host. |
| Edit | Allows you to view or edit the network attachment parameters such as interfaces that you want to attach to the selected network.<br><br>To edit the network attachment information, check the check box next to the network name that you want to edit and choose the **Edit** action. In the **Edit Network Attachment** window, edit the required values, attach or detach the network attachment, click the **Edit** link to edit the CLI freeform config for the switch, and click **Save** to apply the changes or click **Cancel** to discard the changes. The edited network attachment is shown in the table on the **Network Attachments** horizontal tab of the **Networks** tab in the **Fabric Overview** window. |
| Preview | Allows you to preview the configuration of the network attachments for the selected network.<br><br>**Note**<br>This action is not allowed for attachments that are in deployed or NA status.<br><br>To preview the network, check the check box next to the network name and choose **Preview** from **Actions** drop-down list. The **Preview Configuration** window for the fabric appears.<br><br>You can preview the network attachment details such as the network name, fabric name, switch name, serial number, IP address, and role, network status, pending configuration, and progress of the configuration. Click the lines link in the **Pending Config** column to view the lines for which the configuration is pending. Click **Close**. |

| Action Item | Description |
|---|---|
| Deploy | Allows you to deploy the pending configuration of the network attachments, for example, interfaces, for the selected network.<br>**Note**<br>This action is not allowed for attachments that are in deployed or NA status.<br><br>To deploy a network, check the check box next to the network name and choose **Deploy** from **Actions** drop-down list. The **Deploy Configuration** window for the fabric appears.<br><br>You can view the details such as the network name, fabric name, switch name, serial number, IP address, and role, network status, pending configuration, and progress of the configuration. Click the lines link in the **Pending Config** column to view the lines for which the configuration is pending. Click the **Deploy** button. The status and progress of the deployment is displayed in the **Network Status** and **Progress** columns. After the deployment is completed successfully, close the window. |
| Import | Allows you to import information about network attachments for the selected fabric.<br><br>To import the network attachments information, choose **Import**. Browse the directory and select the `.csv` file that contains the network attachments information. Click **Open** and then click **OK**. The network information is imported and displayed in the **Network Attachments** horizontal tab on the **Networks** tab in the **Fabric Overview** window. |
| Export | Allows you to export the information about network attachments to a `.csv` file. The exported file contains information pertaining to each network, including the fabric it belongs to, whether the LAN is attached, the associated VLAN, serial number, interfaces, and freeform configuration details that you saved for network attachments.<br><br>To export network attachments information, choose the **Export** action. Select a location on your local system directory to store the network information and click **Save**. The network information file is exported to your local directory. The file name is appended with the date and time at which the file was exported. |

| Action Item | Description |
|---|---|
| Quick Attach | Allows you to immediately attach an attachment to the selected network. You can select multiple entries and attach them to a network at the same instance.<br>**Note**<br>Interfaces cannot be attached to a network using this action.<br>To quickly attach any attachment to a network, choose **Quick Attach** from **Actions** drop-down list. A message appears to inform that the attach action was successful. |
| Quick Detach | Allows you to immediately detach the selected network from an attachment, for example, a fabric. You can select multiple entries and detach them from an attachment at the same instance.<br>To quickly detach any attachment to a network, choose **Quick Detach** from **Actions** drop-down list. A message appears to inform that the detach action was successful.<br>After quick detach, the switch status is not computed when there is no deploy. Post deploy, the configuration compliance calls at entity level (interface or overlay). |

# Provisioning VXLAN EVPN Fabric with eBGP Underlay

This procedure describes how to create an eBGP VXLAN EVPN with eBGP-based underlay and deploy fabric underlay and overlay eBGP policies. IPv6 underlay is not supported with eBGP EVPN.

## Creating VXLAN EVPN Fabric with eBGP-based Underlay

This procedure shows how to create a new VXLAN EVPN fabric with eBGP based underlay.

1. Choose **LAN** > **Fabrics**.
2. From the **Actions** drop-down list, choose **Create Fabric**.

   The **Create Fabric** window appears. The fields are explained.

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

   **Fabric Template** – Click on this to choose the **BGP Fabric** fabric template. Click **Select**. The fabric settings for creating a standalone fabric appear.
3. The **General Parameters** tab is displayed by default. The fields in this tab are:

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

   **BGP AS Mode** – Choose **Multi-AS** or **Same-Tier-AS**.

   In a **Multi-AS** fabric – Unique AS number per leaf/border.

   In a **Same-Tier-AS** fabric – All leafs share one unique AS and all borders share another unique AS.

In both **Multi-AS** and **Same-Tier-AS**, all the spines in a fabric share one unique AS number. The fabric is identified by the spine switch AS number.

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

**Manual Underlay IP Address Allocation** – Check the check box to disable dynamic underlay IP address allocations.

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

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

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

**Enable Performance Monitoring** – Check the check box to enable performance monitoring.



**Note** Performance Monitoring is supported on switches with NX-OS Release 9.3.6 and later.

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

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

   You can convert a routed fabric to a VXLAN enabled fabric by selecting this option. When the fabric is VXLAN enabled, you can create and deploy overlay networks or VRFs. the procedure for creating and deploying networks or VRFs is the same as in **Data Center VXLAN EVPN**. For more information, see Creating Network for Standalone Fabrics, on page 214 and Creating VRF, on page 205.

   **Routed Fabric** – You must uncheck the **Enable EVPN VXLAN Overlay** check box for Routed fabric (an IP fabric with no VXLAN encapsulation) creation. In a Routed Fabric, you can create and deploy networks. For more information, see Overview of Networks in a Routed Fabric, on page 656.

   Whether you create an eBGP Routed or eBGP VXLAN fabric, the fabric uses eBGP as the control plane to build intra-fabric connectivity. Links between spine and leaf switches are autoconfigured with point-to-point (P2P) numbered IP addresses with eBGP peering built on top.

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

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



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

**Anycast Gateway MAC** – Specifies the anycast gateway MAC address for the leaf switches.

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

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



**Note** VXLAN OAM feature in Cisco NDFC is supported on a single fabric or site only. VXLAN OAM is not supported with Multi-site fabrics.

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

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

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

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

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

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

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

**RP mode** – Choose from the two supported multicast modes of replication, **ASM** (for Any-Source Multicast [ASM]) or **BiDir** (for Bidirectional PIM [BIDIR-PIM]). When you enable multicast mode, only the fields pertaining to that multicast mode is enabled and the fields related to other multicast mode is disabled.



**Note** BIDIR-PIM is supported on Cisco's Cloud Scale Family platforms 9300-EX and 9300-FX/FX2, and with NX-OS Release 9.2(1) and later.

**Underlay RP Loopback ID** – Specifies the loopback ID used for the Rendezvous Point (RP). The default is 254.

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

**Underlay Primary RP Loopback ID** – Specifies the primary loopback ID used for the phantom RP.

**Underlay Backup RP Loopback ID** – Specifies the secondary (or backup) loopback ID used for the Fallback Bidir-PIM phantom RP.

The following Loopback ID options are applicable only when the RP count is 4, if bidir is chosen.

**Underlay Second Backup RP Loopback ID** – Specifies the second backup loopback ID used for the phantom RP.

**Underlay Third Backup RP Loopback ID** – Specifies the third backup loopback ID used for the phantom RP.

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

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

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

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

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

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

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

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

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

**Make vPC Peer Link VLAN as Native VLAN** – Enables vPC peer link VLAN as Native VLAN.

**vPC Peer Keep Alive option** – From the drop-down list, select **management** or **loopback**. To use IP addresses assigned to the management port and the management VRF, select **management**. To use IP addresses assigned to loopback interfaces (in non-management VRF), select **loopback**. If you use IPv6 addresses, you must use loopback IDs.

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

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

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

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

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

**vPC Domain Id** – Specifies the vPC domain ID to be used on all vPC pairs. Otherwise unique vPC domain IDs are used (in increment of 1) for each vPC pair.

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



**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 spines.

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

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

**VTEP Loopback Id** – The loopback interface ID is populated as 1 and it is used for VTEP peering purposes.

**BGP Maximum Paths** – Specifies maximum number for BGP routes to be installed for same prefix on the switches for ECMP.

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

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

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



**Note** Plain text passwords are not supported. Log on to the switch, retrieve the encrypted key. Enter the key in the **BGP Authentication Key** field. For more information, refer to Retrieving the Authentication Key , on page 142.

**Enable PIM Hello Authentication** – Enables the PIM hello authentication.

**PIM Hello Authentication Key** – Specifies the PIM hello authentication key.

**Enable BFD** – Check the **Enable BFD** 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.

NDFC supports BFD within a fabric. 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 BFD configurations requires configurations to be deployed via the per switch freeform or per interface freeform policies.

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

```
feature bfd
```

NDFC with BFD enabled, 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 *Cisco Nexus Dashboard Fabric Controller Compatibility Matrix*.

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

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

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

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

For information about how to retrieve the BFD authentication parameters, see Retrieving the Encrypted BFD Authentication Key, on page 143.

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

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

**Layer 2 Host Interface MTU** – Specifies the MTU for the Layer 2 host interface. This value must be an even number.

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

**CoPP Profile** – From the drop-down list, select the appropriate Control Plane Policing (CoPP) profile policy for the fabric. By default, the **strict** is selected.

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

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

**Enable CDP for Bootstrapped Switch** – Check the check box to enable CDP for switches discovered using Bootstrap.

**Enable NX-API** – Check the check box to enable NX-API on HTTPS. This check box is checked by default.

**Enable NX-API on HTTP** – Specifies enabling of NX-API on HTTP. Check **Enable NX-API on HTTP** and **Enable NX-API** check boxes 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 NDFC, 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 both **Enable NX-API** and **Enable NX-API on HTTP** check boxes, applications use HTTP.

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

For more information, see Strict Configuration Compliance, on page 94.

**Enable AAA IP Authorization** – Enables AAA IP authorization (make sure IP Authorization is enabled in the AAA Server).

**Enable NDFC as Trap Host** – Check the check box to enable NDFC as a trap host.

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

**Greenfield Cleanup Option** – Enable the switch cleanup option for greenfield switches without performing a switch reload. This option is typically recommended only for the Data centers with the Cisco Nexus 9000v Switches.

**Enable Default Queuing Policies** – Check the 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 deploy the configuration. 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 necessary configuration to the peer interface freeform block.

Review the actual queuing policies by opening the policy file in the template editor. From Cisco NDFC Web UI, choose **Operations** > **Template**. 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 FEXs. You can change from the **queuing_policy_default_4q_cloudscale** policy to the **queuing_policy_default_8q_cloudscale** policy only when FEXs are offline.

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

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

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

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

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

8. Click the **Manageability** tab. The fields in this tab are:

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

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

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

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

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

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

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

   **AAA Freeform Config** – Specifies the AAA freeform configs.

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

9. Click the **Bootstrap** tab. The fields in this tab are:

   **Enable Bootstrap** – Check the **Enable Bootstrap** 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** – Check the **Local DHCP Server** check box and enter details for the remaining mandatory fields.

**Enable Local DHCP Server** – Check the **Enable Local DHCP Server** check box to enable DHCP service on NDFC and initiate automatic IP address assignment. When you check this check box, the **DHCP Scope Start Address** and **DHCP Scope End Address** fields become editable.

If you do not check this check box, NDFC 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 NDFC 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 Layer2 adjacent (eth1 or out-of-band subnet must be a /64) or Layer3 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. IPs from this scope are allocated to the switches during the POAP bootstrap process.

**Switch Mgmt Default Gateway** – Specifies the default gateway for the DHCP scope.

**Switch Mgmt IP Subnet Prefix** – Specifies the prefix length for DHCP scope.

**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** – Check the check box to include AAA configs from the **Manageability** tab during device bootup.

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

Copy-paste the running-configuration 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 98.

**DHCPv4/DHCPv6 Multi Subnet Scope** – Specifies the field to enter one subnet scope per line. This field is editable after you check the **Enable Local DHCP Server** check box. The format of the scope should be defined as:

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

For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

10. Click **Configuration Backup** tab. The fields in this tab are:

**Hourly Fabric Backup** – Check the **Hourly Fabric Backup** 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. If there is a configuration push in the previous hour, NDFC takes a backup.

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

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

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

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



**Note** Hourly and scheduled backup processes happen only during the next periodic configuration compliance activity, and there can be a delay of up to an hour.

To trigger an immediate backup, do the following:

a. Choose **LAN** > **Topology**.

b. Click within the specific fabric box. The fabric topology screen comes up.

c. From the **Actions** pane at the left part of the screen, click **Re-Sync Fabric**.

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

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

11. Click the **Flow Monitor** tab. The fields in this tab are:

**Enable Netflow** – Check the **Enable Netflow** check box to enable Netflow on VTEPs for this Fabric. By default, Netflow is disabled. On Enable, NetFlow configuration will be applied to all VTEPS that support netflow.



**Note** When Netflow is enabled on the fabric, you can choose not to have netflow on a particular switch by having a dummy **no_netflow** PTI.

If netflow is not enabled at the fabric level, an error message is generated when you enable netflow at the interface, network, or vrf level. For information about Netflow support for Cisco NDFC, refer Netflow Support, on page 149.

In the **Netflow Exporter** area, click **Actions** > **Add** to add one or more Netflow exporters. This exporter is the receiver of the netflow data. The fields on this tab are:

- **Exporter Name** – Specifies the name of the exporter.
- **IP** – Specifies the IP address of the exporter.
- **VRF** – Specifies the VRF over which the exporter is routed.
- **Source Interface** – Enter the source interface name.
- **UDP Port** – Specifies the UDP port over which the netflow data is exported.

Click **Save** to configure the exporter. Click **Cancel** to discard. You can also choose an existing exporter and select **Actions** > **Edit** or **Actions** > **Delete** to perform relevant actions.

In the **Netflow Record** area, click **Actions** > **Add** to add one or more Netflow records. The fields on this screen are:

- **Record Name** – Specifies the name of the record.
- **Record Template** – Specifies the template for the record. Enter one of the record templates names. From Release 12.0.2, the following two record templates are available for use. You can create custom netflow record templates. Custom record templates saved in the template library are available for use here.
  - **netflow_ipv4_record** – to use the IPv4 record template.
  - **netflow_l2_record** – to use the Layer 2 record template.
- **Is Layer2 Record** – Check the **Is Layer2 Record** check box if the record is for Layer2 netflow.

Click **Save** to configure the report. Click **Cancel** to discard. You can also choose an existing record and select **Actions** > **Edit** or **Actions** > **Delete** to perform relevant actions.

In the **Netflow Monitor** area, click **Actions** > **Add** to add one or more Netflow monitors. The fields on this screen are:

- **Monitor Name** – Specifies the name of the monitor.
- **Record Name** – Specifies the name of the record for the monitor.
- **Exporter1 Name** – Specifies the name of the exporter for the netflow monitor.
- **Exporter2 Name** – (optional) Specifies the name of the secondary exporter for the netflow monitor.

The record name and exporters referred to in each netflow monitor must be defined in "**Netflow Record**" and "**Netflow Exporter**".

Click **Save** to configure the flow monitor. Click **Cancel** to discard. You can also choose an existing monitor and select **Actions** > **Edit** or **Actions** > **Delete** to perform relevant actions.

12. Click on the **Fabric** to view summary in the slide-in pane. Click on the Launch icon to view the **Fabric Overview**.

### Guidelines for VXLAN Fabric With eBGP Underlay

- Brownfield migration is not supported for eBGP fabrics.

- You cannot change the leaf switch AS number after it is created and the configuration is deployed. You must delete the **leaf_bgp_asn** policy and execute **Recalculate & Deploy** to remove BGP configuration related to this AS. Then, add the **leaf_bgp_asn policy** with the new AS number.
- The switch between Multi-AS and Same-Tier-AS modes, remove all manually added BGP policies (including **leaf_bgp_asn** on the leaf switch and the ebgp overlay policies), and execute the **Recalculate & Deploy** operation before the mode change.
- You cannot change or delete the leaf swtch leaf_bgp_asn policy if there are ebgp overlay policies present on the device. You need to delete the ebgp overlay policy first, and then delete the **leaf_bgp_asn** policy.
- The supported roles are leaf, spine, and border only. Any role other than leaf, spine, and border is not supported with VXLAN BGP fabric.
- On the border device, VRF-Lite is supported with manual mode. VXLAN Multi-Site is not supported for VXLAN eBGP fabrics.
- TRM is supported with eBGP fabric.

# Adding Switches

Switch can be added to a single fabric at any point in time. To add switches to a fabric and discover existing or new switches, refer to Adding Switches to a Fabric, on page 279.

# Assigning Switch Roles

To assign roles to switches on Nexus Dashboard Fabric Controller refer to Assigning Switch Roles, on page 297.

# Creating vPC Setup

(Optional) Create a vPC setup for a pair of switches in the fabric. Ensure that the switches have the same roles and are connected to each other. For instructions, refer to vPC Fabric Peering, on page 73.

# Deploying Fabric Underlay eBGP Policies

To deploy fabric underlay eBGP policy, you must manually add the **leaf_bgp_asn** policy on each leaf switch to specify the BGP AS number used on the switch. Implementing the **Recalculate & Deploy** operation afterwards will generate eBGP peering over the physical interface between the leaf and spine switches to exchange underlay reachability information. If **Same-Tier-AS mode** is used, you can deploy the **leaf_bgp_asn** policy on all leafs at the same time as they share the same BGP ASN.

To add a policy to the required switch, see Adding a Policy, on page 195.

# Deploying Fabric Overlay eBGP Policies

You must manually add the eBGP overlay policy for overlay peering. NDFC provides the built-in eBGP leaf and spine overlay peering policy templates that you must manually add to the eBGP leaf and spine switches to form the EVPN overlay peering.

## Deploying Spine Switch Overlay Policies

Add the **ebgp_overlay_spine_all_neighbor** policy on the spine switches. This policy can be deployed on all spine switches at once, since they share the same field values.

The fields on the screen are:

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

**Leaf BGP ASN** – The BGP AS numbers of the leaf switches.

**BGP Update-Source Interface** – This is the source interface for BGP updates. **Underlay Routing Loopback** (default is loopback0) is used for this purpose.

**Enable Tenant Routed Multicast** – (Optional) Check the **Enable Tenant Routed Multicast** check box to enable TRM for handling overlay multicast traffic. TRM enabling must match the fabric setting.

**Enable BGP Authentication** – Check the **Enable BGP Authentication** check box to enable BGP authentication.

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

## Deploying Leaf Switch Overlay Policies

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

The fields on the screen are:

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

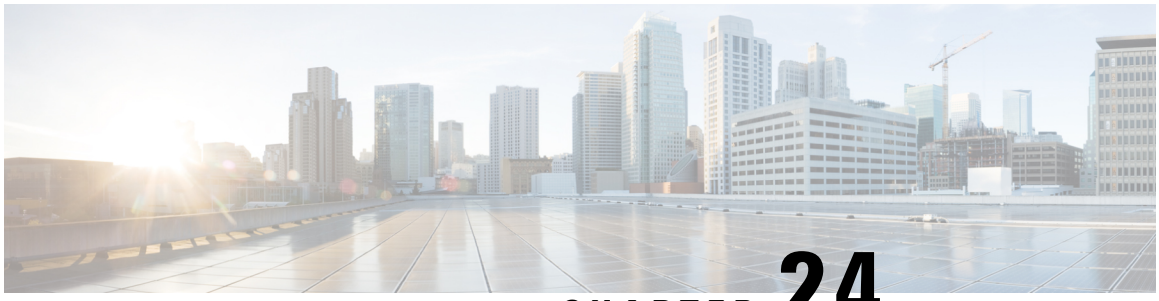
**BGP Update-Source Interface** – This is the source interface for BGP updates. **Underlay Routing Loopback** (default is loopback0) is used for this purpose.

**Enable Tenant Routed Multicast** – (Optional) Check the **Enable Tenant Routed Multicast** check box to enable TRM for handling overlay multicast traffic. TRM enabling must match the fabric setting.

**Enable BGP Authentication** – Check the **Enable BGP Authentication** check box to enable BGP authentication.

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

Click **Actions** > **Recalculate & Deploy**. After the configuration deployment is completed in the **Deploy Configuration** window, click **Close**. You can use the **View/Edit Policy** option to select the policy and click **Push Configuration** to deploy the configuration.



CHAPTER 24

# Managing a Brownfield VXLAN BGP EVPN Fabric


- Overview, on page 551
- Prerequisites, on page 551
- Guidelines and Limitations, on page 552
- Fabric Topology Overview, on page 553
- NDFC Brownfield Deployment Tasks, on page 554
- Verifying the Existing VXLAN BGP EVPN Fabric, on page 554
- Creating a VXLAN EVPN Fabric Using the **Easy_Fabric** Template, on page 557
- Adding Switches and Transitioning VXLAN Fabric Management to NDFC, on page 576
- Configuration Profiles Support for Brownfield Migration, on page 580
- Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration , on page 581
- Migrating a VXLAN EVPN Multi-Site Fabric with Border Gateway Switches , on page 581


## Overview

This use case shows how to migrate an existing VXLAN BGP EVPN fabric to Cisco NDFC. The transition involves migrating existing network configurations to Nexus Dashboard Fabric Controller.

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

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

## Prerequisites

- NDFC-supported NX-OS software versions. For details, refer Cisco Nexus Dashboard Fabric Controller Release Notes.
- Underlay routing protocol is OSPF or IS-IS.
- The following fabric-wide loopback interface IDs must not overlap:
  - Routing loopback interface for IGP/BGP.
  - VTEP loopback ID
  - Underlay rendezvous point loopback ID if ASM is used for multicast replication.

- BGP configuration uses the ‘router-id’, which is the IP address of the routing loopback interface.
- If the iBGP peer template is configured, then it must be configured on the leaf switches and route reflectors. The template name that needs to be used between leaf and route reflector should be identical.
- The BGP route reflector and multicast rendezvous point (if applicable) functions are implemented on spine switches. Leaf switches do not support the functions.
- Familiarity with VXLAN BGP EVPN fabric concepts and functioning of the fabric from the Nexus Dashboard Fabric Controller perspective.
- Fabric switch nodes are operationally stable and functional and all fabric links are up.
- vPC switches and the peer links are up before the migration. Ensure that no configuration updates are in progress or changes pending.
- Create an inventory list of the switches in the fabric with their IP addresses and credentials. Nexus Dashboard Fabric Controller uses this information to connect to the switches.
- Shut down any other controller software you are using presently so that no further configuration changes are made to the VXLAN fabric. Alternatively, disconnect the network interfaces from the controller software (if any) so that no changes are allowed on the switches.
- The switch overlay configurations must have the mandatory configurations defined in the shipping NDFC Universal Overlay profiles. Additional network or VRF overlay related configurations found on the switches are preserved in the freeform configuration associated with the network or VRF NDFC entries.
- All the overlay network and VRF profile parameters such as VLAN name and route map name should be consistent across all devices in the fabric for the brownfield migration to be successful.

# Guidelines and Limitations

- Brownfield import must be completed for the entire fabric by adding all the switches to the NDFC fabric.
- Configuring `terminal color` or variations of this configuration (such as `terminal color persist`) on switches can cause a brownfield import to fail.
- On the **Create Fabric** window, the **Advanced** > **Overlay Mode** fabric setting decides how the overlays will be migrated. If the default config-profile is set, then the VRF and Network overlay configuration profiles will be deployed to switches as part of the migration process. In addition, there will be diffs to remove some of the redundant overlay CLI configurations. These are non network impacting.
- From the **Overlay Mode** drop-down list, if CLI is set, then VRF and Network overlay configurations stay on the switch as-is with no or little changes to address any consistency differences.
- The brownfield import in NDFC supports the simplified NX-OS VXLAN EVPN configuration CLIs. For more information, see Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide, Release 10.2(x).
- The following features are unsupported.
  - Super Spine roles
  - ToR
  - eBGP underlay

    - Layer 3 port channel

- Take a backup of the switch configurations and save them before migration.
- No configuration changes (unless instructed to do so in this document) must be made to the switches until the migration is completed. Else, significant network issues can occur.
- Migration to Cisco Nexus Dashboard Fabric Controller is only supported for Cisco Nexus 9000 switches.
- The Border Spine and Border Gateway Spine roles are supported for the brownfield migration.
- First, note the guidelines for updating the settings. Then update each VXLAN fabric settings as explained below:
    - Some values (BGP AS Number, OSPF, etc) are considered as reference points to your existing fabric, and the values you enter must match the existing fabric values.
    - For some fields (such as IP address range, VXLAN ID range), the values that are auto-populated or entered in the settings are only used for future allocation. The existing fabric values are honored during migration.
    - Some fields relate to new functions that may not exist in your existing fabric (such as advertise-pip). Enable or disable it as per your need.
    - At a later point in time, after the fabric transition is complete, you can update settings if needed.

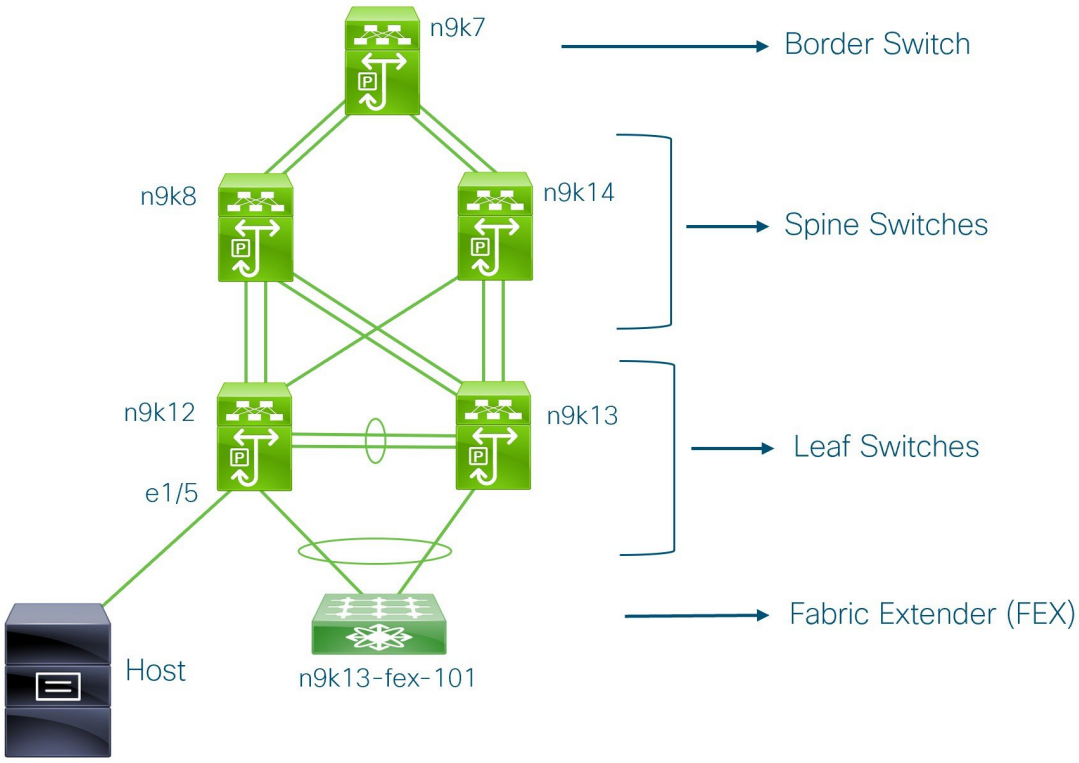
# Fabric Topology Overview

This example use case uses the following hardware and software components:

- Five Cisco Nexus 9000 Series Switches
- One Fabric Extender or FEX
- One host

For information about the supported software images, see *Compatibility Matrix for Cisco NDFC*.

Before we start the transition of the existing fabric, let us see its topology.




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

A host is connected to the n9k12 leaf switch through the interface Ethernet 1/5.

# NDFC Brownfield Deployment Tasks

The following tasks are involved in a Brownfield migration:

1. Verifying the Existing VXLAN BGP EVPN Fabric, on page 554
2. Creating a VXLAN EVPN Fabric Using the Easy_Fabric Template, on page 49
3. Adding Switches and Transitioning VXLAN Fabric Management to NDFC, on page 576

# Verifying the Existing VXLAN BGP EVPN Fabric

Let us check the network connectivity of the **n9k12** switch from the console terminal.

### Procedure

**Step 1** Verify the Network Virtual Interface or NVE in the fabric.

```
n9k12# show nve vni summary
Codes: CP - Control Plane        DP - Data Plane
       UC - Unconfigured

Total CP VNIs: 84    [Up: 84, Down: 0]
Total DP VNIs: 0    [Up: 0, Down: 0]
```

There are 84 VNIs in the control plane and they are up. Before the Brownfield migration, make sure that all the VNIs are up.

**Step 2** Check consistency and failures of vPC.

```
n9k12# show vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 2
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 40
Peer Gateway                      : Enabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled, timer is off.(timeout = 300s)
Delay-restore status              : Timer is off.(timeout = 60s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled
.
.
.
```

**Step 3** Check the EVPN neighbors of the **n9k-12** switch.

```
n9k12# show bgp l2vpn evpn summary
BGP summary information for VRF default, address family L2VPN EVPN
BGP router identifier 192.168.0.4, local AS number 65000
BGP table version is 637, L2VPN EVPN config peers 2, capable peers 2
243 network entries and 318 paths using 57348 bytes of memory
BGP attribute entries [234/37440], BGP AS path entries [0/0]
BGP community entries [0/0], BGP clusterlist entries [2/8]

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.0.0     4 65000     250      91      637    0    0 01:26:59 75
192.168.0.1     4 65000     221      63      637    0    0 00:57:22 75
```

You can see that there are two neighbors corresponding to the spine switches.

Note that the ASN is 65000.

**Step 4** Verify the VRF information.

```
n9k12# show run vrf internet

!Command: show running-config vrf Internet
```

```
!Running configuration last done at: Fri Aug  9 01:38:02 2019
!Time: Fri Aug  9 02:48:03 2019

version 7.0(3)I7(6) Bios:version 07.59

interface Vlan347
  vrf member Internet

interface Vlan349
  vrf member Internet

interface Vlan3962
  vrf member Internet

interface Ethernet1/25
  vrf member Internet

interface Ethernet1/26
  vrf member Internet
vrf context Internet
  description Internet
  vni 16777210
  ip route 204.90.141.0/24 204.90.140.129 name LC-Networks
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
router ospf 300
  vrf Internet
    router-id 204.90.140.3
    redistribute direct route-map allow
    redistribute static route-map static-to-ospf
router bgp 65000
  vrf Internet
    address-family ipv4 unicast
      advertise l2vpn evpn
```

The VRF **Internet** is configured on this switch.

The host connected to the **n9k-12** switch is part of the VRF **Internet**.

You can see the VLANs associated with this VRF.

Specifically, the host is part of **Vlan349**.

**Step 5** Verify the layer 3 interface information.

```
n9k12# show run interface vlan349

!Command: show running-config interface Vlan349
!Running configuration last done at: Fri Aug  9 01:38:02 2019
!Time: Fri Aug  9 02:49:27 2019

version 7.0(3)I7(6) Bios:version 07.59

interface Vlan349
  no shutdown
  vrf member Internet
  no ip redirects
  ip address 204.90.140.134/29
  no ipv6 redirects
  fabric forwarding mode anycast-gateway
```

Note that the IP address is **204.90.140.134**. This IP address is configured as the anycast gateway IP.

**Step 6** Verify the physical interface information. This switch is connected to the Host through the interface Ethernet 1/5.

```
n9k12# show run interface ethernet1/5

!Command: show running-config interface Ethernet1/5
!Running configuration last done at: Fri Aug  9 01:38:02 2019
!Time: Fri Aug  9 02:50:05 2019

version 7.0(3)I7(6) Bios:version 07.59

interface Ethernet1/5
  description to host
  switchport mode trunk
  switchport trunk native vlan 349
  switchport trunk allowed vlan 349,800,815
  spanning-tree bpduguard enable
  mtu 9050
```

You can see that this interface is connected to the host and is configured with VLAN 349.

**Step 7** Verify the connectivity from the host to the anycast gateway IP address.

```
host# ping 204.90.140.134 count unlimited interval 1
PING 204.90.140.134 (204.90.140.134): 56 data bytes
64 bytes from 204.90.140.134: icmp_seq=0 ttl=254 time=1.078 ms
64 bytes from 204.90.140.134: icmp_seq=1 ttl=254 time=1.129 ms
64 bytes from 204.90.140.134: icmp_seq=2 ttl=254 time=1.151 ms
64 bytes from 204.90.140.134: icmp_seq=3 ttl=254 time=1.162 ms
64 bytes from 204.90.140.134: icmp_seq=4 ttl=254 time=1.84 ms
64 bytes from 204.90.140.134: icmp_seq=5 ttl=254 time=1.258 ms
64 bytes from 204.90.140.134: icmp_seq=6 ttl=254 time=1.273 ms
64 bytes from 204.90.140.134: icmp_seq=7 ttl=254 time=1.143 ms
```

We let the ping command run in the background while we transition the existing brownfield fabric into Nexus Dashboard Fabric Controller.

# Creating a VXLAN EVPN Fabric Using the Easy_Fabric Template

This topic describes how to create a new VXLAN EVPN fabric using the **Easy_Fabric** template and contains descriptions for the IPv4 underlay. For information about the IPv6 underlay, see IPv6 Underlay Support for Easy Fabric, on page 69.

1. Navigate to the **LAN Fabrics** page:

   **LAN** > **Fabrics**

2. Click **Actions** > **Create Fabric**.

   The **Create Fabric** window appears.

3. Enter a unique name for the fabric in the **Fabric Name** field, then click **Choose Fabric**.

   A list of all available fabric templates are listed.

4. From the available list of fabric templates, choose the **Easy_Fabric** template, then click **Select**.

5. Enter the necessary field values to create a fabric.

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



**Note** If you're 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), see Multi-Site Domain for VXLAN BGP EVPN Fabrics , on page 589 before creating the member fabric.

- General Parameters, on page 558
- Replication, on page 560
- VPC, on page 561
- Protocols, on page 562
- Advanced, on page 565
- Resources, on page 569
- Manageability, on page 572
- Bootstrap, on page 573
- Configuration Backup, on page 574
- Flow Monitor, on page 575

6. When you have completed the necessary configurations, click **Save**.
   - Click on the fabric to display a summary in the slide-in pane.
   - Click on the Launch icon to display the Fabric Overview.

### General Parameters

The **General Parameters** tab is displayed by default. The fields in this tab are described in the following table.

| Field | Description |
|---|---|
| **BGP ASN** | Enter the BGP AS number the fabric is associated with. This must be same as existing fabric. |
| **Enable IPv6 Underlay** | Enable the IPv6 underlay feature. For information, see IPv6 Underlay Support for Easy Fabric, on page 69. |
| **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. |

| Field | Description |
|---|---|
| **Underlay Subnet IPv6 Mask** | Specifies the subnet mask for the fabric interface IPv6 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.<br><br>To deploy spine devices as RRs, Nexus Dashboard Fabric Controller 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 won't change.<br><br>*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.<br><br>*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.<br><br>**1.** Change the value in the drop-down box to 2.<br><br>**2.** Identify the spine switches designated as route reflectors.<br><br>An instance of the **rr_state** policy is applied on the spine switch if it's 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.<br><br>**3.** Delete the unneeded spine devices from the fabric (right-click the spine switch icon and choose **Discovery > Remove from fabric**).<br><br>If you delete existing RR devices, the next available spine switch is selected as the replacement RR.<br><br>**4.** Click **Deploy Config** in the fabric topology window.<br><br>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. |
| **Enable Performance Monitoring** | Check the check box to enable performance monitoring.<br><br>Ensure that you do not clear interface counters from the Command Line Interface of the switches. Clearing interface counters can cause the Performance Monitor to display incorrect data for traffic utilization. If you must clear the counters and the switch has both `clear counters` and `clear counters snmp` commands (not all switches have the `clear counters snmp` command), ensure that you run both the main and the SNMP commands simultaneously. For example, you must run the `clear counters interface ethernet slot/port` command followed by the `clear counters interface ethernet slot/port snmp` command. This can lead to a one time spike. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Replication

The fields in the **Replication** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **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.<br>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.<br>The replication mode change isn't allowed if a policy template instance is created for the current mode. For example, if a multicast related policy is created and deployed, you can't change the mode to Ingress. |
| **Enable Tenant Routed Multicast (TRM)** | Check 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**.<br>For more information, see Overview of Tenant Routed Multicast, on page 69. |
| **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]).<br>When you choose ASM, the BiDir related fields aren't enabled. When you choose BiDir, the BiDir related fields are enabled.<br>**Note**<br>BIDIR-PIM is supported on Cisco's Cloud Scale Family platforms 9300-EX and 9300-FX/FX2, and software release 9.2(1) onwards.<br>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. |
| **Underlay Primary RP Loopback ID** | Enabled if you choose BIDIR-PIM as the multicast mode of replication.<br>The primary loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay. |
| **Underlay Backup RP Loopback ID** | Enabled if you choose BIDIR-PIM as the multicast mode of replication.<br>The secondary loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay. |
| **Underlay Second Backup RP Loopback Id** | Used for the second fallback Bidir-PIM Phantom RP. |

| Field | Description |
|---|---|
| **Underlay Third Backup RP Loopback Id** | Used for the third fallback Bidir-PIM Phantom RP. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### VPC

The fields in the **VPC** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **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.<br>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. Uncheck the check box to disable the function. |
| **vPC advertise-pip** | Select the check box to enable the Advertise PIP feature.<br>You can enable the advertise PIP feature on a specific vPC as well. . |
| **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. .<br>**Note**<br>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**. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Protocols

The fields in the **Protocols** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **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 Anycast Loopback Id** | The loopback interface ID is greyed out and used for vPC Peering in VXLANv6 Fabrics only. |
| **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.<br>**Note**<br>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.<br>**Note**<br>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 Network Point-to-Point** | Enables network point-to-point on fabric interfaces which are numbered. |
| **Enable IS-IS Authentication** | Select the check box to enable IS-IS authentication. Deselect the check box to disable it. If you enable this field, the IS-IS authentication fields are enabled. |
| **IS-IS Authentication Keychain Name** | Enter the Keychain name, such as CiscoisisAuth. |
| **IS-IS Authentication Key ID** | The Key ID is populated. |

| Field | Description |
|---|---|
| **IS-IS Authentication Key** | Enter the Cisco Type 7 encrypted key.<br>**Note**<br>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. |
| **Set IS-IS Overload Bit** | When enabled, set the overload bit for an elapsed time after a reload. |
| **IS-IS Overload Bit Elapsed Time** | Allows you to clear the overload bit after an elapsed time in seconds. |
| **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.<br>**Note**<br>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.<br>**Note**<br>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** | Select this check box to enable PIM hello authentication on all the intra-fabric interfaces of the switches in a fabric. This check box is editable only for the Multicast replication mode. Note this check box is valid only for the IPv4 underlay. |
| **PIM Hello Authentication Key** | Specifies the PIM hello authentication key. For more information, see Retrieving PIM Hello Authentication Key.<br>To retrieve the PIM Hello Authentication Key, perform the following steps:<br>**1.** SSH into the switch.<br>**2.** On an unused switch interface, enable the following:<br>`switch(config)# interface e1/32`<br>`switch(config-if)# ip pim hello-authentication ah-md5 pimHelloPassword`<br>In this example, **pimHelloPassword** is the cleartext password that has been used.<br>**3.** Enter the **show run interface** command to retrieve the PIM hello authentication key.<br>`switch(config-if)# show run interface e1/32 \| grep pim`<br>`ip pim sparse-mode`<br>`ip pim hello-authentication ah-md5 3 d34e6c5abc7fecf1caa3b588b09078e0`<br>In this example, **d34e6c5abc7fecf1caa3b588b09078e0** is the PIM hello authentication key that should be specified in the fabric settings. |

| Field | Description |
|---|---|
| **Enable BFD** | Check 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.<br>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.<br>The following config is pushed after you select the **Enable BFD** check box:<br>`feature bfd`<br>For information about BFD feature compatibility, refer your respective platform documentation and for information about the supported software images, see *Compatibility Matrix for Cisco Nexus Dashboard Fabric Controller*. |
| **Enable BFD for iBGP** | Check the check box to enable BFD for the iBGP neighbor. This option is disabled by default. |
| **Enable BFD for OSPF** | Check 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** | Check 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** | Check 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.<br>Following are examples of the BFD global policies:<br>`router ospf <ospf tag>`<br>`  bfd`<br>`router isis <isis tag>`<br>`  address-family ipv4 unicast`<br>`  bfd`<br>`ip pim bfd`<br>`router bgp <bgp asn>`<br>`  neighbor <neighbor ip>`<br>`  bfd` |
| **Enable BFD Authentication** | Check the check box to enable BFD authentication. If you enable this field, the **BFD Authentication Key ID** and **BFD Authentication Key** fields are editable.<br>**Note**<br>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.<br>For information about how to retrieve the BFD authentication parameters. . |

| Field | Description |
|---|---|
| **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.<br>If you use BGP templates, add the authentication configuration within the template and uncheck the Enable BGP Authentication check box to avoid duplicate configuration.<br>In the sample configuration, the 3DES password is displayed after password 3.<br>`router bgp 65000`<br>`  password 3 sd8478fswerdfw3434fsw4f4w34sdsd8478fswerdfw3434fsw4f4w`<br>The following fields can be used to specify different configurations:<br>• **iBGP Peer-Template Config** – Specifies the config used for RR and spines with border role.<br>• **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).<br>In a 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. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Advanced

The fields in the **Advanced** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **VRF Template** | Specifies the VRF template for creating VRFs. |
| **Network Template** | Specifies the network template for creating networks. |
| **VRF Extension Template** | Specifies the VRF extension template for enabling VRF extension to other fabrics. |
| **Network Extension Template** | Specifies the network extension template for extending networks to other fabrics. |
| **Overlay Mode** | VRF/Network configuration using config-profile or CLI, default is config-profile. For more information, see Overlay Mode, on page 83. |
| **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. |

| Field | Description |
|---|---|
| **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. |
| **Unshut Host Interfaces by Default** | Check this check box to unshut the host interfaces by default. |
| **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.<br><br>The following list describes the variables in the syntax:<br>• **$$VNI$$**: Specifies the network VNI ID found in the switch configuration. This is a mandatory keyword required to create unique network names.<br>• **$$VLAN_ID$$**: Specifies the VLAN ID associated with the network.<br>VLAN ID is specific to switches, hence Nexus Dashboard Fabric Controller picks the VLAN ID from one of the switches, where the network is found, randomly and use it in the name.<br>We recommend not to use this unless the VLAN ID is consistent across the fabric for the VNI.<br>• **<string>**: This variable is optional and you can enter any number of alphanumeric characters that meet the network name guidelines.<br><br>An example overlay network name: Site_VNI12345_VLAN1234<br><br>**Note**<br>Ignore this field for greenfield deployments. The **Brownfield Overlay Network Name Format** applies for the following brownfield imports:<br>• CLI-based overlays<br>• Configuration profile-based overlay |
| **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. |

| Field | Description |
|---|---|
| **Enable VXLAN OAM** | Enables the VXLAM OAM functionality for devices in the fabric. This is enabled by default. Uncheck the check box to disable VXLAN OAM function.<br>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.<br>**Note**<br>The VXLAN OAM feature in Cisco Nexus Dashboard Fabric Controller is only supported on a single fabric or site. |
| **Enable Tenant DHCP** | Check 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.<br>**Note**<br>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 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 Nexus Dashboard Fabric Controller, 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.<br>**Note**<br>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)** | Check 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. It enables bi-directional compliance checks to flag additional configs in the running config that are not in the intent/expected config. By default, this feature is disabled. |
| **Enable AAA IP Authorization** | Enables AAA IP authorization, when IP Authorization is enabled in the remote authentication server. This is required to support Nexus Dashboard Fabric Controller in scenarios where customers have strict control of which IP addresses can have access to the switches. |
| **Enable NDFC as Trap Host** | Select this check box to enable Nexus Dashboard Fabric Controller as an SNMP trap destination. Typically, for a native HA Nexus Dashboard Fabric Controller deployment, the eth1 VIP IP address will be configured as SNMP trap destination on the switches. By default, this check box is enabled. |
| **Anycast Border Gateway advertise-pip** | Enables to advertise Anycast Border Gateway PIP as VTEP. Effective on MSD fabric 'Recalculate Config'. |

| Field | Description |
|---|---|
| **Greenfield Cleanup Option** | Enable the switch cleanup option for switches imported into Nexus Dashboard Fabric Controller 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 check 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 79. |
| **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 Nexus Dashboard Fabric Controller.<br>If the PTP loopback ID is not found during **Deploy Config**, the following error is generated:<br>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** | Check the check box to enable the MPLS Handoff feature. For more information, see the MPLS SR and LDP Handoff, on page 619 chapter in External/WAN Layer 3 Connectivity for VXLAN BGP EVPN Fabrics. |
| **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. 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.<br>Review the actual queuing policies by opening the policy file in the template editor. From Cisco Nexus Dashboard Fabric Controller Web UI, choose **Operations > Templates**. Search for the queuing policies by the policy file name, for example, **queuing_policy_default_8q_cloudscale**. Choose the file. From the **Actions** drop-down list, select **Edit template content** to edit the policy.<br>See the *Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide* for platform specific details. |

| Field | Description |
|---|---|
| **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 Enabling MACsec.<br>*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 should be added via the switch freeform on NDFC. For more information, see Enabling Freeform Configurations on Fabric Switches , on page 95. |
| **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. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Resources

The fields in the **Resources** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **Manual Underlay IP Address Allocation** | *Do not* check this check box if you are transitioning your VXLAN fabric management to Nexus Dashboard Fabric Controller.<br>• By default, Nexus Dashboard Fabric Controller allocates the underlay IP address resources (for loopbacks, fabric interfaces, etc) dynamically from the defined pools. If you check the check box, the allocation scheme switches to static, and some of the dynamic IP address range fields are disabled.<br>• For static allocation, the underlay IP address resources must be populated into the Resource Manager (RM) using REST APIs.<br>• The Underlay RP Loopback IP Range field stays enabled if BIDIR-PIM function is chosen for multicast replication.<br>• 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.<br>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. |
| **Underlay Routing Loopback IPv6 Range** | Specifies Loopback0 IPv6 Address Range |
| **Underlay VTEP Loopback IPv6 Range** | Specifies Loopback1 and Anycast Loopback IPv6 Address Range. |
| **Underlay Subnet IPv6 Range** | Specifies IPv6 Address range to assign Numbered and Peer Link SVI IPs. |
| **BGP Router ID Range for IPv6 Underlay** | Specifies BGP router ID range for IPv6 underlay. |
| **Layer 2 VXLAN VNI Range** | Specifies the overlay VXLAN VNI range for the fabric (min:1, max:16777214). |
| **Layer 3 VXLAN VNI Range** | Specifies the overlay VRF VNI range for the fabric (min:1, max:16777214). |
| **Network VLAN Range** | VLAN range for the per switch overlay network (min:2, max:4094). |

| Field | Description |
|---|---|
| **VRF VLAN Range** | VLAN range for the per switch overlay Layer 3 VRF (min:2, max:4094). |
| **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.<br>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 Back2Back&ToExternal, then VRF Lite IFCs are auto-created. |
| **Auto Deploy for Peer** | This check box is applicable for VRF Lite deployment. When you select this checkbox, auto-created VRF Lite IFCs will have the **Auto Generate Configuration for Peer** field in the **VRF Lite** tab set.<br>To access VRF Lite IFC configuration, navigate to the **Links** tab, select the particular link, and then choose **Actions** > **Edit**.<br>You can check or uncheck the check box when the **VRF Lite Deployment** field is not set to **Manual**. This configuration only affects the new auto-created IFCs and does not affect the existing IFCs. You can edit an auto-created IFC and check or uncheck the **Auto Generate Configuration for Peer** field. This setting takes priority always. |
| **Auto Deploy Default VRF** | When you select this check box, the **Auto Generate Configuration on default VRF** field is automatically enabled for auto-created VRF Lite IFCs. You can check or uncheck this check box when the **VRF Lite Deployment** field is not set to **Manual**. The **Auto Generate Configuration on default VRF** field when set, automatically configures the physical interface for the border device, and establishes an EBGP connection between the border device and the edge device or another border device in a different VXLAN EVPN fabric. |
| **Auto Deploy Default VRF for Peer** | When you select this check box, the **Auto Generate Configuration for NX-OS Peer on default VRF** field is automatically enabled for auto-created VRF Lite IFCs. You can check or uncheck this check box when the **VRF Lite Deployment** field is not set to **Manual**. The **Auto Generate Configuration for NX-OS Peer on default VRF** field when set, automatically configures the physical interface and the EBGP commands for the peer NX-OS switch.<br>**Note**<br>To access the **Auto Generate Configuration on default VRF** and **Auto Generate Configuration for NX-OS Peer on default VRF** fields for an IFC link, navigate to the **Links** tab, select the particular link and choose **Actions** > **Edit**. |
| **Redistribute BGP Route-map Name** | Defines the route map for redistributing the BGP routes in default VRF. |

| Field | Description |
|---|---|
| **VRF Lite Subnet IP Range** and **VRF Lite Subnet Mask** | These fields are populated with the DCI subnet details. Update the fields as needed.<br>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:<br>**Note**<br>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.<br>**1.** Update the L2 range and click **Save**.<br>**2.** 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. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Manageability

The fields in the **Manageability** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **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.<br>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. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Bootstrap

The fields in the **Bootstrap** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **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.<br>After you enable bootstrap, you can enable the DHCP server for automatic IP address assignment using one of the following methods:<br>• External DHCP Server: Enter information about the external DHCP server in the **Switch Mgmt Default Gateway** and **Switch Mgmt IP Subnet Prefix** fields.<br>• 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.<br>If you do not select this check box, Nexus Dashboard Fabric Controller 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.<br>**Note**<br>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.<br>*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. |

| Field | Description |
|---|---|
| **Enable AAA Config** | Select this check box to include AAA configurations from the Manageability tab as part of the device start-up config post bootstrap. |
| **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.<br>The format of the scope should be defined as:<br>**DHCP Scope Start Address, DHCP Scope End Address, Switch Management Default Gateway, Switch Management Subnet Prefix**<br>For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24 |
| **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.<br>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 Enabling Freeform Configurations on Fabric Switches , on page 95. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Configuration Backup

The fields in the **Configuration Backup** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **Hourly Fabric Backup** | Select the check box to enable an hourly backup of fabric configurations and the intent.<br>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. |

| Field | Description |
|---|---|
| **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.<br><br>Select both the check boxes to enable both back up processes.<br><br>The backup process is initiated after you click **Save**.<br><br>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.<br><br>The number of fabric backups that will be retained on NDFC is decided by the **Settings** > **Server Settings** > **LAN Fabric** > **Maximum Backups per Fabric**.<br><br>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.<br><br>**Note**<br>To trigger an immediate backup, do the following:<br><br>**1.** Choose **LAN > Topology**.<br><br>**2.** Click within the specific fabric box. The fabric topology screen comes up.<br><br>**3.** From the **Actions** pane at the left part of the screen, click **Re-Sync Fabric**.<br><br>You can also initiate the fabric backup in the fabric topology window. Click **Backup Now** in the **Actions** pane. |

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

### Flow Monitor

The fields in the **Flow Monitor** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

| Field | Description |
|---|---|
| **Enable Netflow** | Check this check box to enable Netflow on VTEPs for this Fabric. By default, Netflow is disabled. On Enable, NetFlow configuration will be applied to all VTEPS that support netflow.<br><br>**Note**<br>When Netflow is enabled on the fabric, you can choose not to have netflow on a particular switch by having a dummy no_netflow PTI.<br><br>If netflow is not enabled at the fabric level, an error message is generated when you enable netflow at the interface, network, or vrf level. For information about Netflow support for Cisco NDFC, refer to Netflow Support, on page 149. |

In the **Netflow Exporter** area, click **Actions > Add** to add one or more Netflow exporters. This exporter is the receiver of the netflow data. The fields on this screen are:

- **Exporter Name** – Specifies the name of the exporter.

- **IP** – Specifies the IP address of the exporter.
- **VRF** – Specifies the VRF over which the exporter is routed.
- **Source Interface** – Enter the source interface name.
- **UDP Port** – Specifies the UDP port over which the netflow data is exported.

Click **Save** to configure the exporter. Click **Cancel** to discard. You can also choose an existing exporter and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

In the **Netflow Record** area, click **Actions > Add** to add one or more Netflow records. The fields on this screen are:

- **Record Name** – Specifies the name of the record.
- **Record Template** – Specifies the template for the record. Enter one of the record templates names. In Release 12.0.2, the following two record templates are available for use. You can create custom netflow record templates. Custom record templates saved in the template library are available for use here.
  - **netflow_ipv4_record** – to use the IPv4 record template.
  - **netflow_l2_record** – to use the Layer 2 record template.
- **Is Layer2 Record** – Check this check box if the record is for Layer2 netflow.

Click **Save** to configure the report. Click **Cancel** to discard. You can also choose an existing record and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

In the **Netflow Monitor** area, click **Actions > Add** to add one or more Netflow monitors. The fields on this screen are:

- **Monitor Name** – Specifies the name of the monitor.
- **Record Name** – Specifies the name of the record for the monitor.
- **Exporter1 Name** – Specifies the name of the exporter for the netflow monitor.
- **Exporter2 Name** – (optional) Specifies the name of the secondary exporter for the netflow monitor.

The record name and exporters referred to in each netflow monitor must be defined in "**Netflow Record**" and "**Netflow Exporter**".

Click **Save** to configure the monitor. Click **Cancel** to discard. You can also choose an existing monitor and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

**What's next:** Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

# Adding Switches and Transitioning VXLAN Fabric Management to NDFC

Let us discover and add switches to the newly created fabric.

### Procedure

**Step 1** Double click on the newly created fabric to view the **Fabric Overview** screen.

Click on **Switches** tab.

**Step 2** From the **Actions** drop-down list, select **Add Switches**.

The **Add Switches** window appears.

Similarly, you can add switches on **Topology** window. On Topology window, choose a fabric, right-click on a fabric and click **Add Switches**.

**Step 3** On the **Add Switches - Fabric** screen, enter the **Seed Switch Details**.

Enter the IP address of the switch in the **Seed IP** field. Enter the username and password of the switches that you want to discover.

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

Make sure to check the **Preserve Config** check box. This ensures that the current configuration of the switches will be retained.

**Step 4** Click **Discover Switches**.

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

**Step 5** Check the check box next to the switches that have to be imported into the fabric and click **Import into fabric**.

It is best practice to discover multiple switches at the same time in a single attempt. The switches must be cabled and connected to the NDFC server and the switch status must be manageable.

If switches are imported in multiple attempts, then please ensure that all the switches are added to the fabric before proceeding with the Brownfield import process.

**Step 6** Click **Import into fabric**.

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

**Note**
You should not close the screen and try to import switches again until all selected switches are imported or an error message comes up.

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

**Step 7** After a successful import, the progress bar shows **Done** for all the switches. Click **Close**.

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

At this point, you must not try to add Greenfield or *new* switches. Support is not available for adding new switches during the migration process. It might lead to undesirable consequences for your network. However, you can add a new switch after the migration process is complete.

**Step 8** After all the network elements are discovered, they are displayed in the **Topology** window in a connected topology. Each switch is assigned the **Leaf** role by default.

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

You should not proceed to the next step until all switches in the existing fabric are discovered in NDFC.

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

**Note**
The supported roles for switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images are Border Leaf, Border Spine, Leaf, and Spine

**Step 9** Select the switch, click **Actions** > **Set Role**. On the Select Role screen, select **Border** and click **Select**.

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

**Note**
You need to manually create a vPC pairing when the L3 keep alive is configured on the switches. Otherwise, the vPC configuration is automatically picked up from the switches.

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

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

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

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

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

**Step 10** From the Fabric Overview **Actions** drop-down list, choose **Recalculate and Deploy**.

When you click **Recalculate and Deploy**, NDFC obtains switch configurations and populates the state of every switch from the current running configuration to the current expected configuration, which is the intended state maintained in NDFC.

If there are configuration mismatches, **Pending Config** column shows the number of lines of difference. Click on the Pending Config column to view the **Pending Config** and **Side-by-Side Comparison** of the running configuration. Click **Deploy** to apply the configurations.

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

**Note**

- The brownfield migration requires best practices to be followed on the existing fabric such as maintain consistency of the overlay configurations.
- The Brownfield migration may take some time to complete since it involves collecting the running configuration from switches, build the NDFC configuration intent based on these, consistency checks etc.
- Any errors or inconsistencies that are found during the migration is reported in fabric errors. The switches continue to remain in the Migration mode. You should fix these errors and complete the migration again by clicking **Deploy** until no errors are reported.

**Step 11** After the configurations are generated, review them by clicking the links in the **Preview Config** column.

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

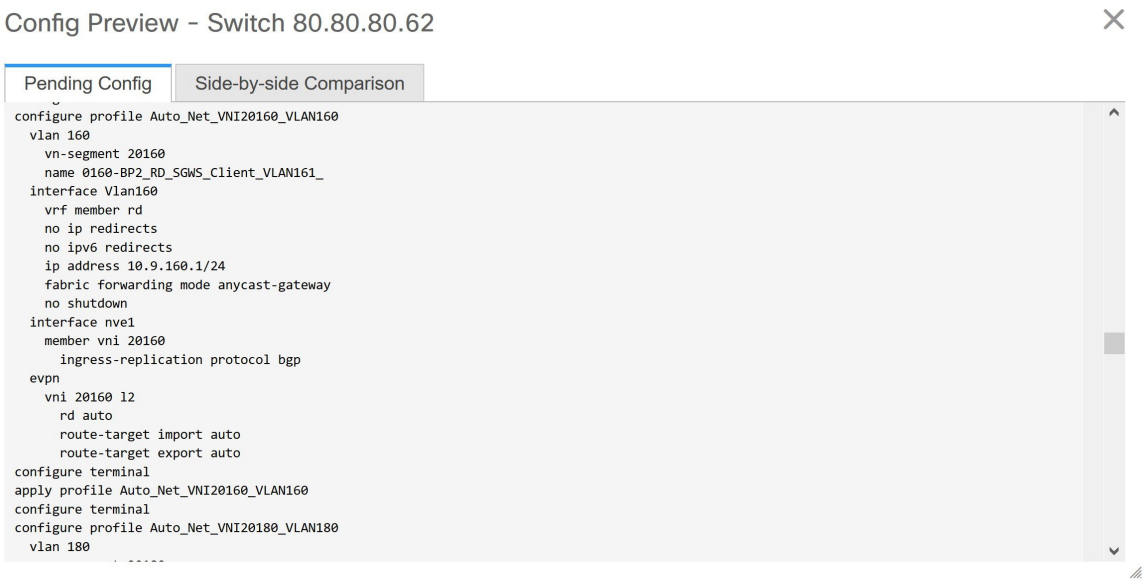
The Side-by-side Comparison tab displays the running configuration and expected configuration side-by-side.

The **Pending Config** tab displays the set of configurations that need to be deployed on a switch in order to go from the current running configuration to the current expected or intended configuration.

The **Pending Config** tab may show many configuration lines that will be deployed to the switches. Typically, on a successful brownfield import, these lines correspond to the configuration profiles pushed to the switches for a overlay network configuration. Note that the existing network and VRF-related overlay configurations are not removed from the switches.

**Note**

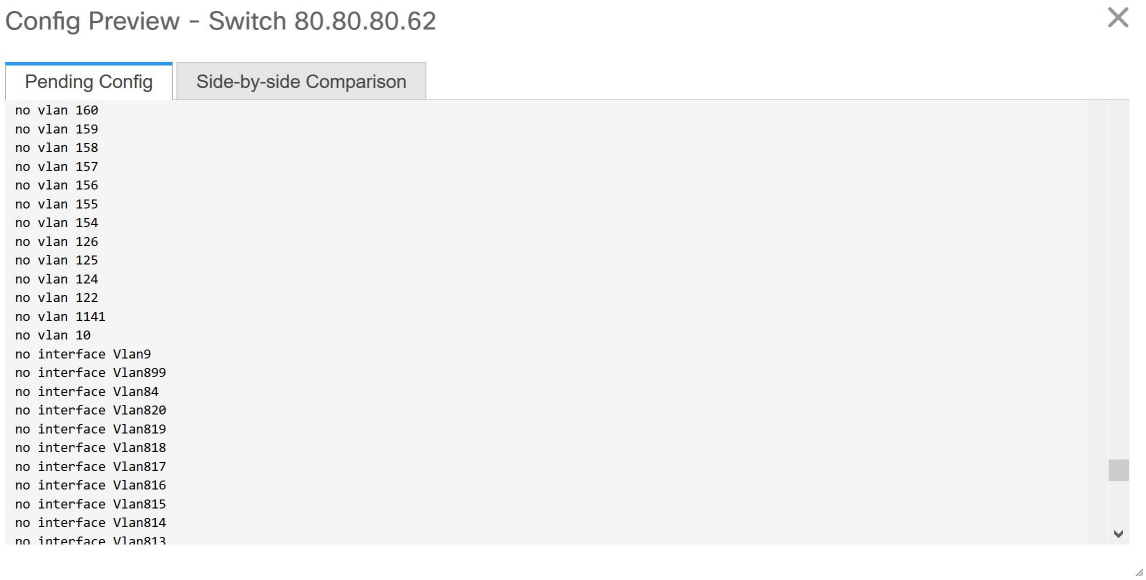
The configuration profiles are NDFC required constructs for managing the VXLAN configurations on the switches. During the Brownfield import process, they capture the same information as the original VXLAN configurations already present on the switches. In the following image, the configuration profile with **vlan 160** is applied.




As part of the import process, after the configuration profiles are applied, the original CLI based configuration references will be removed from the switches. These are the ‘no’ CLIs that will be seen towards the end of

the diffs. The VXLAN configurations on the switches will be persisted in the configuration profiles. In the following image, you can see that the configurations will be removed, specifically, **no vlan 160**.

The removal of CLI based configuration is allowed if the **Overlay Mode** is set to **config-profile**, and not CLI.




The **Side-by-side Comparison** tab displays the Running Config and Expected Config on the switch.

**Step 12** Close the **Config Preview Switch** window after reviewing the configurations.

**Step 13** Click **Deploy Config** to deploy the pending configuration onto the switches.

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

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

In the fabric topology screen that comes up, all imported switch instances are displayed in green color, indicating successful configuration. Also, the **Migration Mode** label is not displayed on any switch icon.

NDFC has successfully imported a VXLAN-EVPN fabric.

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

For more information, see Fabric Overview, on page 181.

# Configuration Profiles Support for Brownfield Migration

Cisco NDFC supports the Brownfield import of fabrics with VXLAN overlay provisioned with configuration profiles. This import process recreates the overlay configuration intent based on the configuration profiles. The underlay migration is performed with the usual Brownfield migration.

This feature can be used to recover your existing Easy fabric when a NDFC backup is not available to be restored. In this case, you must install the latest NDFC release, create a fabric, and then import the switches into the fabric.

Note that this feature is not recommended for the NDFC upgrade. For more information, see *Cisco NDFC Installation and Upgrade Guide*.

The following are the guidelines for the support of configuration profiles:

- The Brownfield migration of configuration profiles is supported for the **Data Center VXLAN EVPN** template.
- The configuration profiles on the switches must be a subset of the default overlay **Universal** profiles. If extra configuration lines are present that are not part of the **Universal** profiles, unwanted profile refreshes will be seen. In this case, after you recaluclate and deploy configuration, review the diffs using the **Side-by-side Comparison** feature and deploy the changes.
- Brownfield migration with switches having a combination of VXLAN overlay configuration profiles and regular CLIs is not supported. If this condition is detected, an error is generated, and migration is aborted. All the overlays must be with either configuration profiles or regular CLIs only.

# Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration

After brownfield migration, if you add new spine or leaf switches, you should manually configure the PIM-BIDIR feature.

The following procedure shows how to manually configure the PIM-BIDIR feature for a new Leaf or Spine:

**Procedure**

**Step 1** Check the **base_pim_bidir_11_1** policies that are created for an RP added through the brownfield migration. Check the RP IP and Multicast Group used in each **ip pim rp-address** *RP_IP* **group-list** *MULTICAST_GROUP* **bidir** command.

**Step 2** Add respective **base_pim_bidir_11_1** policies from the **View/Edit Policies** window for the new Leaf or Spine, push the config for each **base_pim_bidir_11_1** policy.

# Migrating a VXLAN EVPN Multi-Site Fabric with Border Gateway Switches

When you migrate an existing VXLAN EVPN Multi-Site fabric with a border gateway switch into NDFC, make sure to note the following guidelines:

- Uncheck all **Auto** IFC creation related fabric settings. Review the settings and ensure they are unchecked as follows:

- **Data Center VXLAN EVPN** fabric

  Uncheck **Auto Deploy Both** check box under **Resources** tab.

- **VXLAN EVPN Multi-Site** fabric

  Uncheck **Multi-Site Underlay IFC Auto Deployment Flag** check box under **DCI** tab.

- Underlay Multisite peering: The eBGP peering and corresponding routed interfaces for underlay extensions between sites are captured in **switch_freeform** and **routed_inerfaces**, and optionally in the **interface_freeform** configs. This configuration includes all the global configs for multisite. Loopbacks for EVPN multisite are also captured via the appropriate interface templates.
- Overlay Multisite peering: The eBGP peering is captured as part of **switch_freeform** as the only relevant config is under **router bgp**.
- Overlays containing Networks or VRFs: The corresponding intent is captured with the profiles on the Border Gateways with **extension_type = MULTISITE**.

1. Create all the required fabrics including the Data Center VXLAN EVPN and Multi-Site Interconnect Network fabrics with the required fabric settings. Disable the Auto VRF-Lite options as mentioned above. For more information, refer to *Creating VXLAN EVPN Fabric* and *External Fabric* sections.
2. Import all the switches into all the required fabrics and set roles accordingly.
3. Click **Recalculate and Deploy** in each of the fabrics and ensure that the Brownfield Migration process reaches the 'Deployment' phase. Now, do not click **Deploy Configuration**.
4. Create the **VXLAN EVPN Multi-Site** fabric with the required fabric settings and disable the **Auto MultiSite IFC** options as shown in Guidelines. For more information, see *Creating a VXLAN EVPN Multi-Site Fabric*.
5. Move all the member fabrics into the VXLAN EVPN Multi-Site fabric. Do not proceed further till this step is completed successfully. For more information, see *Moving the Member1 Fabric Under VXLAN EVPN Multi-Site-Parent-Fabric*.



**Note** The Overlay Networks and VRFs definitions in each of the Easy Fabrics must be symmetric for them to get added successfully to the VXLAN EVPN Multi-Site fabric. Errors will be reported if any mismatches are found. These must be fixed by updating the overlay information in the fabric(s) and added to the VXLAN EVPN Multi-Site fabric.

6. Create all the Multisite Underlay IFCs such that they match the IP address and settings of the deployed configuration.



**Note** Additional interface configurations must be added to the Source/Destination interface freeform fields in the **Advanced** section as needed.

For more information, see *Configuring Multi-Site Overlay IFCs*.

7. Create all the Multisite Overlay IFCs such that they match the IP address and settings of the deployed configuration. You will need to add the IFC links. For more information, see *Configuring Multi-Site Overlay IFCs*.
8. If there are VRF-Lite IFCs also, create them as well.



**Note** If the Brownfield Migration is for the case where Configuration Profiles already exist on the switches, the VRF-Lite IFCs will be created automatically in Step #3.

9. If Tenant Routed Multicast (TRM) is enabled in the VXLAN EVPN Multi-Site fabric, edit all the TRM related VRFs and Network entries in VXLAN EVPN Multi-Site and enable the TRM parameters.

   This step needs to be performed if TRM is enabled in the fabric. If TRM is not enabled, you still need to edit each Network entry and save it.
10. Now click **Recalculate and Deploy** in the VXLAN EVPN Multi-Site fabric, but, do not click **Deploy Configuration**.
11. Navigate to each member fabric, click **Recalculate and Deploy**, and then click **Deploy Configuration**.

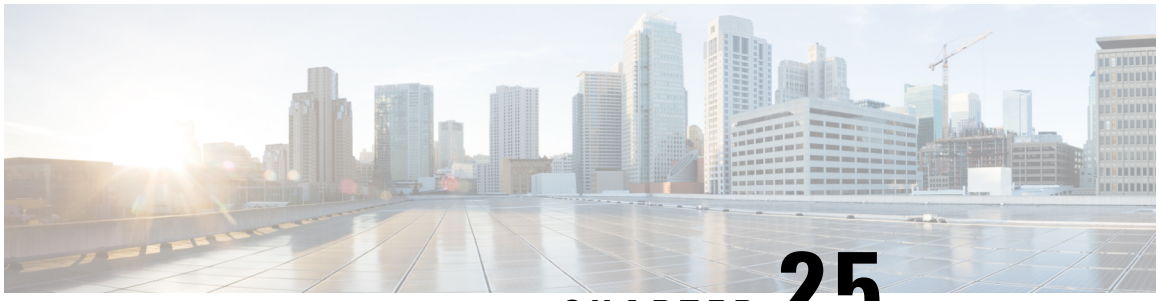
This completes the Brownfield Migration. You can now manage all the networks or VRFs for BGWs by using the regular NDFC Overlay workflows.

When you migrate an existing VXLAN EVPN Multi-Site fabric with border gateway switches (BGW) that has a Layer-3 port-channel for Underlay IFCs, make sure to do the following steps:



**Note** Ensure that the child fabrics are added into VXLAN EVPN Multi-Site before migrating the VXLAN EVPN Multi-Site fabric.

1. Click on appropriate VXLAN EVPN Multi-Site child fabric and navigate to **Fabrics** > **Interfaces** to view the BGW. Choose an appropriate Layer-3 port channel to use for underlay IFC.
2. On **Policy** column, choose **int_port_channel_trunk_host_11_1** from drop-down list. Enter the associated port-channel interface members and then click **Save**.
3. Navigate to the **Tabular view** of the VXLAN EVPN Multi-Site fabric. Edit layer-3 port link, choose the multisite underlay IFC link template, enter source and destination IP addresses. These IP addresses are the same as existing configuration values on the switches
4. Do the steps from step 7 to 11 from above procedure.



CHAPTER 25

# Configuring a VXLANv6 Fabric

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


- Overview, on page 585
- Creating VXLAN EVPN Fabric with IPv6 Underlay, on page 586


## Overview

From Cisco NDFC, you can create an Easy fabric with IPv6 only underlay. The IPv6 underlay is supported only for the **Data Center VXLAN EVPN** template. In the IPv6 underlay fabric, intra-fabric links, routing loopback, vPC peer link SVI, and NVE loopback interface for VTEP are configured with IPv6 addresses. EVPN BGP neighbor peering is also established using IPv6 addressing.

The following guidelines are applicable for IPv6 underlay:

- IPv6 underlay is supported for the Cisco Nexus 9000 Series switches with Cisco NX-OS Release 9.3(1) or higher.
- VXLANv6 is only supported Cisco Nexus 9332C, Cisco Nexus C9364C, and Cisco Nexus modules that end with EX, GX, FX, FX2, FX3, or FXP.

**Note** VXLANv6 is defined as a VXLAN fabric with IPv6 underlay.

- In VXLANv6, the platforms supported on spine are all Nexus 9000 Series and Nexus 3000 Series platforms.
- The overlay routing protocol supported for the IPv6 fabric is BGP EVPN.
- vPC with physical multichassis EtherChannel trunk (MCT) feature is supported for the IPv6 underlay network in NDFC. The vPC peer keep-alive can be loopback or management with IPv4 or IPv6 address.
- Brownfield migration is supported for the VXLANv6 fabrics. Note that L3 vPC keep-alive using IPv6 address is not supported for brownfield migration. This vPC configuration is deleted after the migration. However, L3 vPC keep-alive using IPv4 address is supported.
- DHCPv6 is supported for the IPv6 underlay network.
- The following features are not supported for VXLAN IPv6 underlay:

- Multicast underlay
- Tenant Routed Multicast (TRM)
- ISIS, OSPF, and BGP authentication
- VXLAN Multi-Site
- Dual stack underlay
- vPC Fabric Peering
- DCI SR-MPLS or MPLS-LDP handoff
- BFD
- Super Spine switch roles
- NGOAM

# Creating VXLAN EVPN Fabric with IPv6 Underlay

This procedure shows how to create a VXLAN EVPN fabric with IPv6 underlay. Note that only the fields for creating a VXLAN fabric with IPv6 underlay are documented. For information about the remaining fields, see Creating a VXLAN EVPN Fabric Using the Easy_Fabric Template, on page 49.

**Procedure**

**Step 1** Choose **LAN** > **Fabrics**.

**Step 2** From the **Actions** drop-down list, choose **Create Fabric**.

The **Create Fabric** window appears.

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

**Fabric Template** – From the drop-down list, choose **Data Center VXLAN EVPN**.

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

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

**Enable IPv6 Underlay** – Check the **Enable IPv6 Underlay** check box .

**Enable IPv6 Link-Local Address** – Check the **Enable IPv6 Link-Local Address** check box to use the link local addresses in the fabric between leaf-spine and spine-border interfaces. If you check this check box, the **Underlay Subnet IPv6 Mask** field is not editable. By default, the **Enable IPv6 Link-Local Address** field is enabled.

IPv6 underlay supports **p2p** networks only. Therefore, the **Fabric Interface Numbering** drop-down list is disabled.

**Underlay Subnet IPv6 Mask** – Specify the subnet mask for the fabric interface IPv6 addresses.

**Underlay Routing Protocol** – Specify the IGP used in the fabric, that is, OSPF or IS-IS for VXLANv6.

**Step 4** All the fields under the **Replication** tab are disabled.

IPv6 underlay supports ingress replication mode only.

**Step 5** Click the **VPC** tab.

**vPC Peer Keep Alive option** – Choose **management** or **loopback**. To use IP addresses assigned to the management port and the management VRF, choose management. To use IP addresses assigned to loopback interfaces and a non-management VRF, choose underlay routing loopback with IPv6 address for PKA. Both the options are supported for IPv6 underlay.

**Step 6** Click the **Protocols** tab.

**Underlay Anycast Loopback Id** – Specify the underlay anycast loopback ID for IPv6 underlay. You cannot configure IPv6 address as secondary, an additional loopback interface is allocated on each vPC device. Its IPv6 address is used as the VIP.

**Step 7** Click the **Resources** tab.

**Manual Underlay IP Address Allocation**: Check the check box to manually allocate underlay IP addresses. The dynamic underlay IP addresses fields are disabled.

**Underlay Routing Loopback IPv6 Range**: Specify loopback IPv6 addresses for protocol peering.

**Underlay VTEP Loopback IPv6 Range**: Specify loopback IPv6 addresses for VTEPs.

**Underlay Subnet IPv6 Range**: Specify the IPv6 address range that is used for assigning IP addresses for numbered and peer link SVIs. To edit this field, uncheck **Enable IPv6 Link-Local Address** check box under the **General Parameters** tab.

**BGP Router ID Range for IPv6 Underlay**: Specify the address range to assign BGP Router IDs. The IPv4 addressing is used for router with BGP and underlay routing protocols.

**Step 8** Click the **Bootstrap** tab.

**Enable Bootstrap**: Check the **Enable Bootstrap** check box. If this check box is not chosen, none of the other fields on this tab are editable.

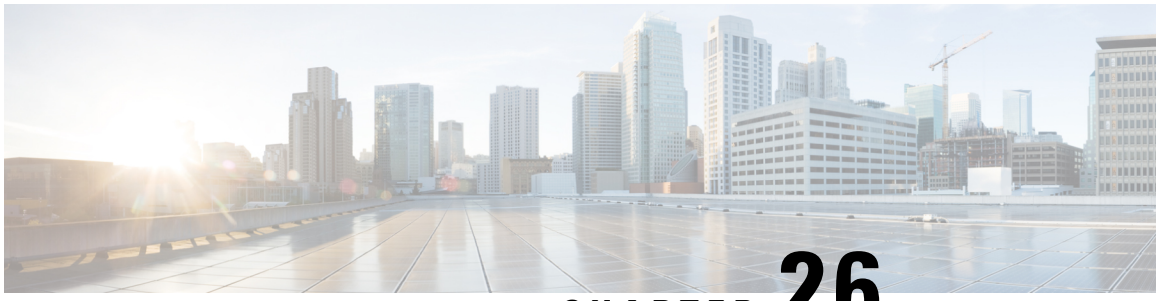
**Enable Local DHCP Server**: Check the check box to initiate automatic assignment of IP addresses assignment through the local DHCP server. The **DHCP Scope Start Address** and **DHCP Scope End Address** fields are editable only after you check this check box.

**DHCP Version**: Choose DHCPv4 from the drop-down list.

**Step 9** Click **Save** to complete the creation of the fabric.

### What to do next

Adding Switches to a Fabric



CHAPTER 26

# Multi-Site Domain for VXLAN BGP EVPN Fabrics


- Multi-Site Domain for VXLAN BGP EVPN Fabrics , on page 589
- MSD and Member Fabric Process Flow, on page 590
- Creating the VXLAN EVPN Multi-Site and Associating Member Fabrics, on page 593
- Creating and Deploying Networks and VRFs in an MSD Fabric, on page 599
- Moving a Standalone Fabric (With Existing Networks and VRFs) to an MSD Fabric , on page 600
- Support for CloudSec in Multi-Site Deployment, on page 601


## 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.

As 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.

The topology view for the MSD fabric displays all member fabrics, and how they are connected to each other, in one view. 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**

- The VXLAN OAM feature in Cisco NDFC is only supported on a single fabric or site.
- After you unpair a BGW vPC, perform a **Recalculate Config** and **Deploy Config** on the member fabric followed by a **Recalculate Config** and **Deploy Config** 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 *Data Center VXLAN EVPN*) 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 networks (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.

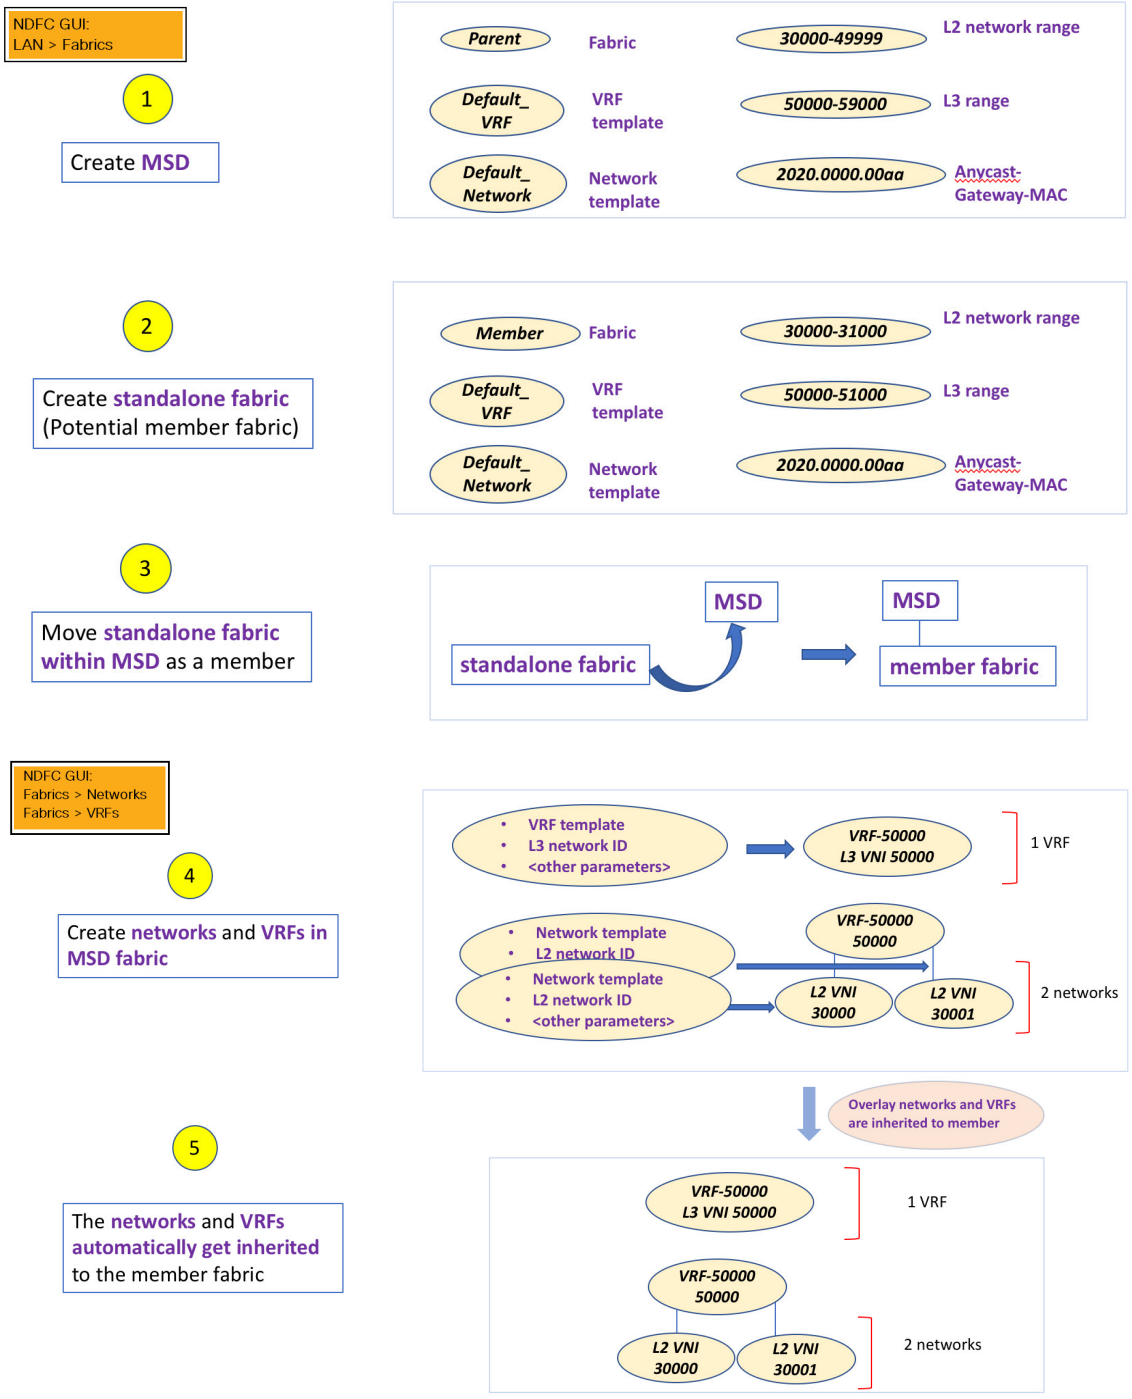
Fabric instance values can only be edited or updated in the fabric context from the VRFs and Networks window. Double click on the appropriate fabric to view **Fabric Overview** and choose **Networks** or **VRFs** tab. Some of the examples of fabric instance variables are BGP ASN, Multicast group per network or VRF, etc. For information about editing multicast group address, see Creating Networks in the MSD Fabric, on page 599.

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

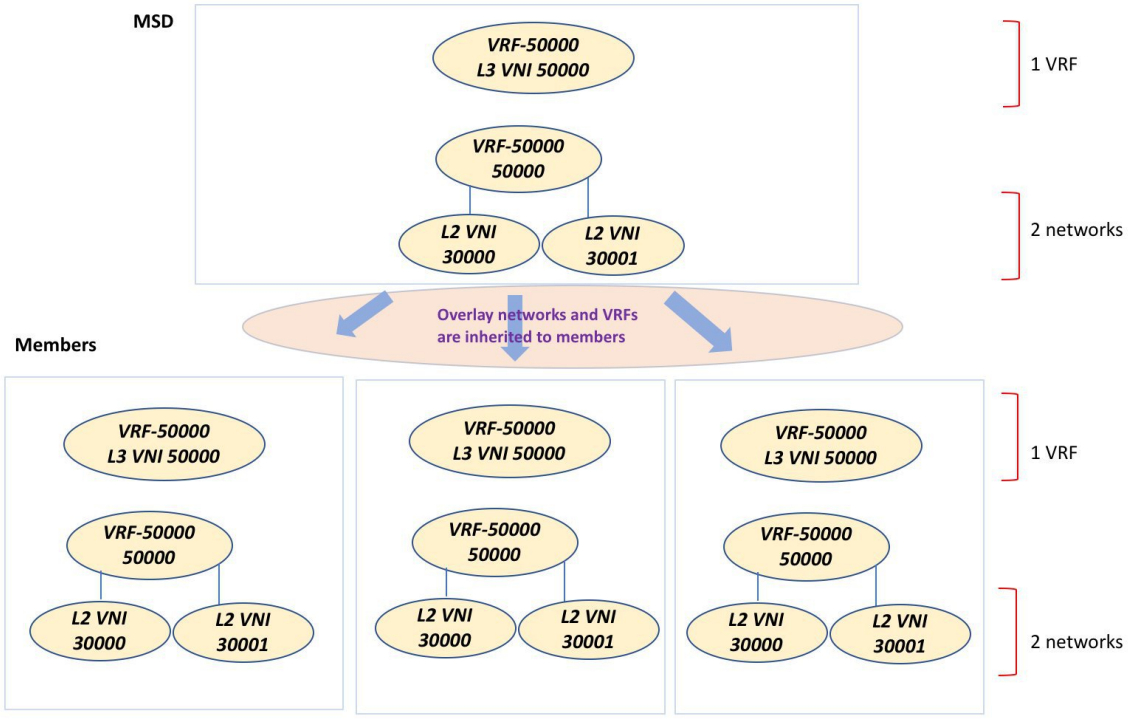
# MSD and Member Fabric Process Flow

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

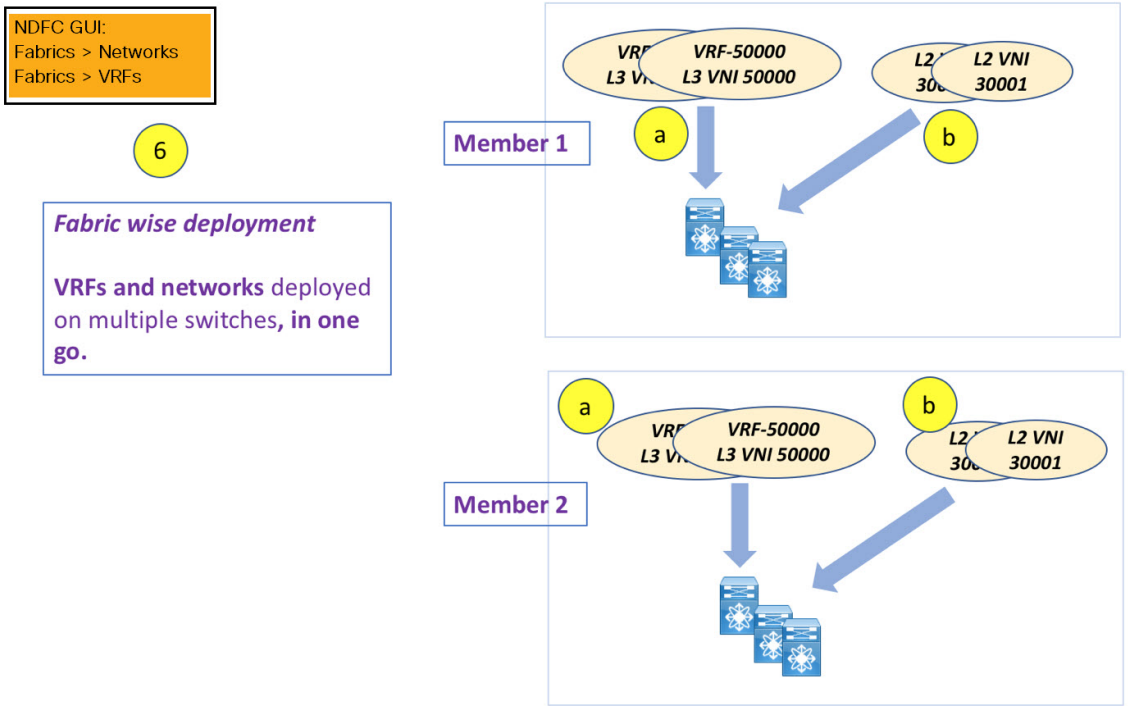
A high-level flow chart of the MSD and member fabric creation and MSD-to-member fabric inheritance process:




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




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



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, NDFC 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 the VXLAN EVPN Multi-Site and Associating Member Fabrics

The process is explained in two steps:

1. Create the VXLAN EVPN Multi-Site fabric.
2. Create a new standalone fabric and move it under the VXLAN EVPN Multi-Site fabric as a member fabric.

### Creating the VXLAN EVPN Multi-Site fabric

1. From **Actions** drop-down list, choose **Create Fabric**.

   The **Create Fabric** window appears.

2. Enter a unique name for the Fabric.

   Click on **Choose Template** to pick a template for your fabric.

   A list of all available fabric templates are listed.

3. From the available list of Fabric templates, choose the **VXLAN EVPN Multi-Site** fabric template.

   Click **Select**.

   Enter the necessary field values to create a Fabric.

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

4. In the **General Parameters** 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 multisite 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 **Recalculate and Deploy** in the VXLAN EVPN Multi-Site fabric.

5. Click the **DCI** tab.

   The fields are:

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

   **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 601.

   **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.

6. Click the **Resources** tab.

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

   A unique loopback IP address is assigned from this range to each member fabric because each member site must have a Multi-site Routing Loopback 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.

7. Click the **Configuration Backup** tab.

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

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

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

   The backup process is initiated after you click Save.

8. Click **Save**.

   A message appears briefly at the bottom right part of the screen, indicating that you have created a new VXLAN EVPN Multi-Site fabric. After fabric creation, the fabric page comes up.

   When a new VXLAN EVPN Multi-Site fabric is created, the newly created VXLAN EVPN Multi-Site fabric instance appears in the Fabrics table.

   The VXLAN EVPN Multi-Site fabric is displayed, and it contains the member fabric names as a branch. When no member fabric is created, it is displayed as a standalone fabric.

The steps for creation of the VXLAN EVPN Multi-Site fabric and moving member fabrics under it are:

1. Create the VXLAN EVPN Multi-Site fabric.
2. **Create a new standalone fabric and move it under the VXLAN EVPN Multi-Site 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 VXLAN EVPN Multi-Site Fabric as a Member

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

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

The values under the **Resources** tab 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 VXLAN EVPN Multi-Site 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 VXLAN EVPN Multi-Site fabric. Else, member fabric movement to the VXLAN EVPN Multi-Site fabric fail.

Other pointers:

- 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 in the list of fabrics.

### Moving the Member1 Fabric Under VXLAN EVPN Multi-Site-Parent-Fabric

You should go to the VXLAN EVPN Multi-Site fabric Overview to associate a member fabric under it.

1. Double click on the VXLAN EVPN Multi-Site fabric to view the **Fabric Overview** screen.
2. On the **Child Fabrics** tab, click .

   You can also click on **Fabric Overview > Actions > Add Child Fabrics** to add member fabrics to the VXLAN EVPN Multi-Site fabric.

   A list of child fabrics that are not part of any VXLAN EVPN Multi-Site fabric appears. Member fabrics of other VXLAN EVPN Multi-Site fabric container fabrics are not displayed here.
3. As *Member1* fabric is to be associated with the VXLAN EVPN Multi-Site fabric, select the **Member1** fabric and click **Select**.
4. Select the Fabric and click **Select**.

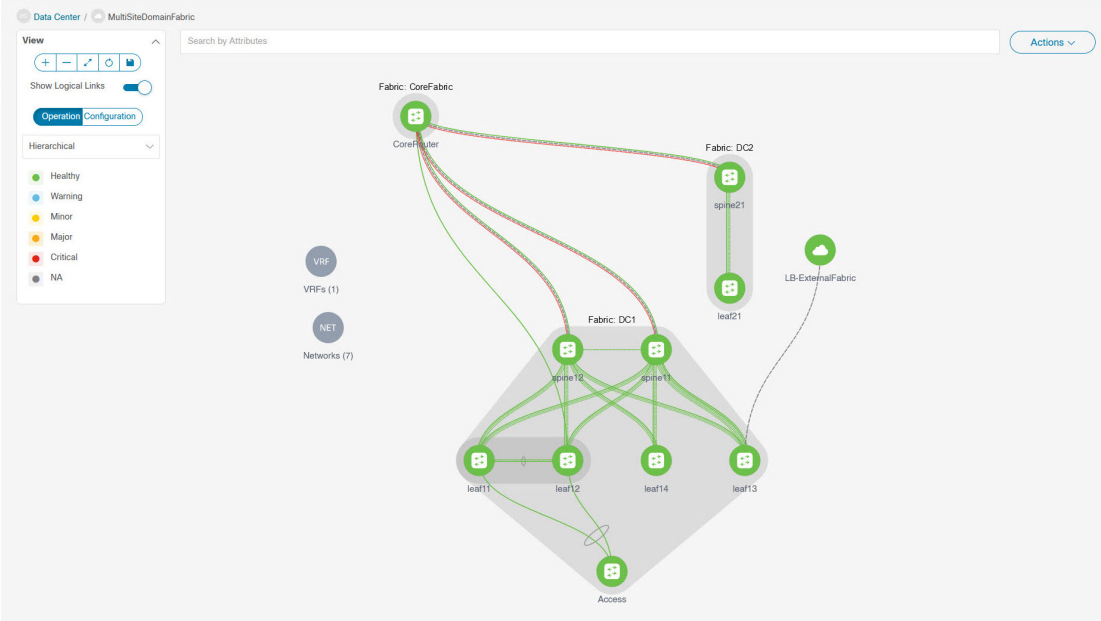
   You can see that *Member1* is now added to VXLAN EVPN Multi-Site fabric and is displayed in the **Child Fabrics** in the Fabrics list table.

### VXLAN EVPN Multi-Site Fabric Topology View Pointers

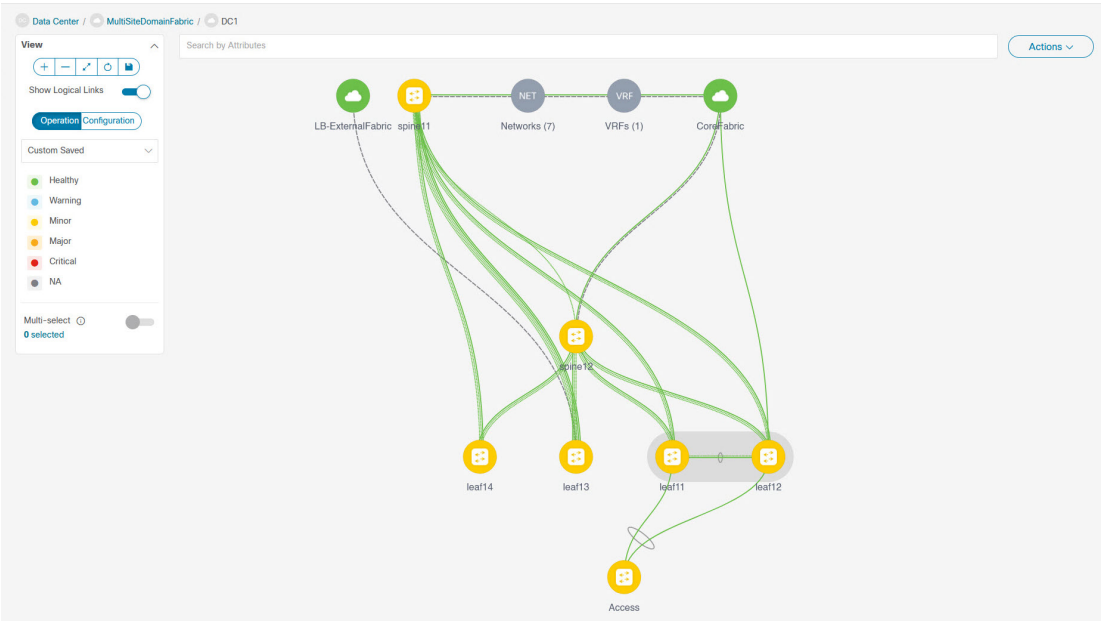
The Topology tab displays the configured VXLAN EVPN Multi-Site fabrics and its child fabrics.

- **VXLAN EVPN Multi-Site fabric topology view** – VXLAN EVPN Multi-Site fabric and their member fabrics displayed. A boundary defines each member fabric. All fabric devices of the fabric are confined to the boundary.

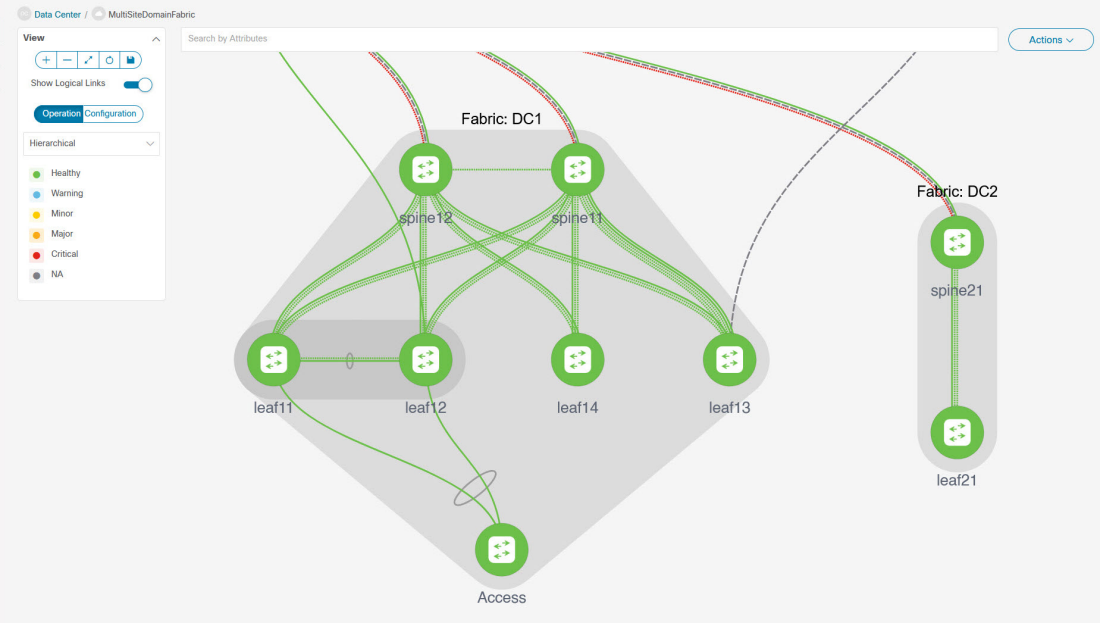
  Double click on the member fabric to view further elements.




- **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 the VXLAN EVPN Multi-Site 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, NDFC 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.




### Configuring External Fabric Routers for VXLAN EVPN Child Fabric Connectivity

To successfully pair external fabric routers with VXLAN EVPN child fabrics using the **Fabric-Group Settings** for auto-deploy configurations, the following prerequisites and configurations must be applied to the external fabric routers:

- **Role of the device:** The device in the external fabric must be set as a core-router.
- **Loopback0 Interface Configuration:** The loopback0 interface must be configured with an IP address and tag:

  ```
  interface loopback0

  ip address <ip>/<mask> tag <tag>
  ```

- **BGP Configuration:**BGP must be configured on the external fabric router:

  ```
  feature bgp

  router bgp <ASN>
  router-id <ip>
  address-family ipv4 unicast
  ```

  You can use the ASN, along with the router ID, when defining the auto-deploy configuration in the **Fabric-Group Settings** within NDFC.

  These configurations are crucial for establishing the necessary control plane and routing capabilities, enabling seamless integration and auto-deployment of IFCs between the external fabric and the VXLAN EVPN child fabrics.

### Adding and Editing Links

To add a link, choose **Actions** > **More** > **Add Link**. To edit a link, choose **Actions** > **More** > **Edit Link**.

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.

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.

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.

### Creating Networks in the MSD Fabric

Some guidelines and pointers:

- In the MSD fabric level, if the **Enable L3 Gateway on Border** check box is selected and you upgrade the NDFC service, then it is automatically removed from the MSD fabric level during upgrade.
- You can edit the **Network Profile** part (**General** and **Advanced** tabs) of the MSD fabric network.
- 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*.
- You can only delete networks from the MSD fabric, and not member fabrics. You must undeploy the networks on the respective fabric devices before deletion.
- When you delete networks from the MSD fabric, the networks are automatically removed from the member fabrics too.

See Creating Networks for the Standalone Fabric.

### Creating VRFs in the MSD 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.

See Creating VRF.

### Deleting Networks and VRFs in the MSD and Member Fabrics

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

1. Undeploy the networks on the respective fabric devices before deletion.
2. Delete the networks from the MSD fabric.
3. Undeploy the VRFs on the respective fabric devices before deletion.
4. Delete the VRFs from the MSD fabric. 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.

# 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. NDFC validates the standalone fabric (network and VRF information) against the (network and VRF information) of the MSD fabric to avoid conflict entries. An example of conflict entries is two common network names with a different network ID. After validation and there is no conflict, the standalone fabric is moved to the MSD fabric as a member fabric. Details:

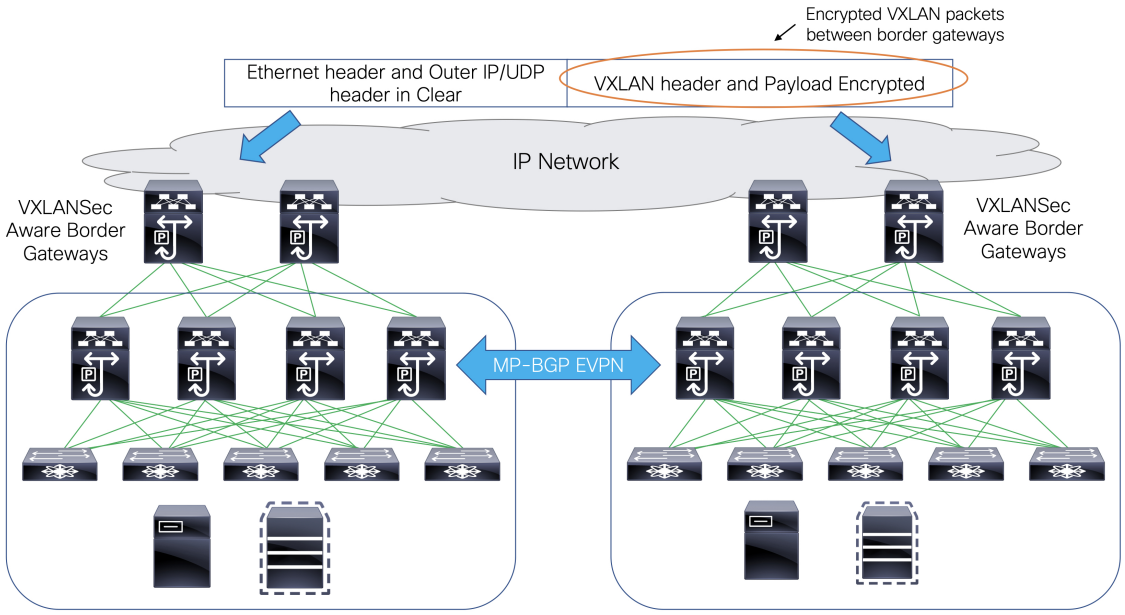
- The MSD fabric inherits the networks and VRFs of the standalone fabric that do not exist in the MSD fabric. These networks and VRFs are in turn inherited by the member fabrics.
- The newly created member fabric inherits the networks and VRFs of the MSD fabric (that do not exist in the newly created member fabric).

- If there are conflicts between the standalone and MSD fabrics, validation ensures that an error message is displayed. You can move the standalone fabric to MSD again after updating. If the move is 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.

# 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.

You can enable CloudSec while creating 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.

## Enabling CloudSec in MSD

On the NDFC Web UI, choose **LAN > Fabircs**. You can either create a new MSD fabric by clicking **Create Fabric** or edit the existing MSD fabric by clicking **Edit Fabric**.

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

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

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

When you click **Recalculate & Deploy**, you can verify theses configs in the **Preview Config** window for the border gateway switches.



**Note** CloudSec isn't supported if the border gateway has vPC or TRM is enabled on it, that is, TRM enabled on multisite overlay IFC. If CloudSec is enabled in this scenario, appropriate warning or error messages are generated.

**CloudSec Key String** – Specifies the hex key string. Enter a 66 hexadecimal string if you choose **AES_128_CMAC** or enter a 130 hexadecimal string if you choose **AES_256_CMAC**.

**CloudSec Cryptographic Algorithm** – Choose **AES_128_CMAC** or **AES_256_CMAC**.

**CloudSec Enforcement** – Specifies whether the CloudSec enforcement should be strict or loose.

**strict** – Deploys the CloudSec configuration to all the border gateways in fabrics in MSD. If there are any border gateways that don't support CloudSec, then an error message is generated, and the configuration isn't pushed to any switch.

If **strict** is chosen, the **tunnel-encryption must-secure** CLI is pushed to the CloudSec enabled gateways within MSD.

**loose** – Deploys the CloudSec configuration to all the border gateways in fabrics in MSD. If there are any border gateways that don't support CloudSec, then a warning message is generated. In this case, the CloudSec config is only deployed to the switches that support CloudSec. If **loose** is chosen, the **tunnel-encryption must-secure** CLI is removed if it exists.



**Note** There should be at least two fabrics in MSD with border gateways that support CloudSec. If there's only one fabric with a CloudSec capable device, then the following error message is generated:

CloudSec needs to have at least 2 sites that can support CloudSec.

To remove this error, meet the criteria of having at least two sites that can support CloudSec or disable CloudSec.

**CloudSec Status Report Timer** – Specifies the CloudSec Operational Status periodic report timer in minutes. This value specifies how often the NDFC 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 NDFC, 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 NDFC to form the key chain policy. NDFC forms the encryption-policy by taking all default values. NDFC pushes the same key chain policy, the same encryption-policy, and encryption-peer policies to each CloudSec capable gateways. On each gateway, there's one encryption-peer policy for each remote gateway that is CloudSec capable, using the same keychain and same key policy.

If you don't want to use the same key for the whole MSD fabric or want to enable CloudSec on a subset of all sites, you can use **switch_freeform** to manually push the CloudSec config to the switches.

Capture all the CloudSec config in **switch_freeform**.

For example, the below config is included in the **switch_freeform** policy:

```
feature tunnel-encryption
evpn multisite border-gateway 600
  dci-advertise-pip
tunnel-encryption must-secure-policy
tunnel-encryption policy CloudSec_Policy1
tunnel-encryption source-interface loopback20
key chain CloudSec_Key_Chain1 tunnel-encryption
  key 1000
    key-octet-string 7 075e731f1a5c4f524f43595f507f7d737062677147524054590 70b0b0701585440
    cryptographic-algorithm AES_128_CMA
tunnel-encryption peer-ip 192.168.0.6
  keychain CloudSec_Key_Chain1 policy CloudSec_Policy1
```

Add **tunnel-encryption** in the Freeform Config of the uplink interface policy which will generate config like the following:

```
interface ethernet1/13
  no switchport
  ip address 192.168.1.14/24 tag 54321
  evpn multisite dci-tracking
  tunnel-encryption
  mtu 9216
  no shutdown
```

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

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 NDFC, 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 NDFC 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

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** Select **Actions** > **Detailed View**.

**Step 3** Click the **Link** tab and choose **CloudSec Operational View** tab on the left.

**Step 4** If CloudSec is disabled, the **CloudSec Operational View** isn't displayed.

The **Operational View** has the following fields and descriptions.

| Fields | Description |
|---|---|
| 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:<br>• Up: The CloudSec session is successfully established between the switches.<br>• 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. |

**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** On the NDFC Web UI, choose **Operations > Programmable Reports**.

**Step 2** Click **Create Report**.

**Step 3** Specify a unique name for the report in the **Report Name** field.

**Step 4** From the **Select Template** drop-down list, select **fabric_cloudsec_oper_status**.

**Step 5** Click **Next** to view the **Source & Recurrence** tab.

**Step 6** In the **Recurrence** field, choose the frequency at which the report job should be run.

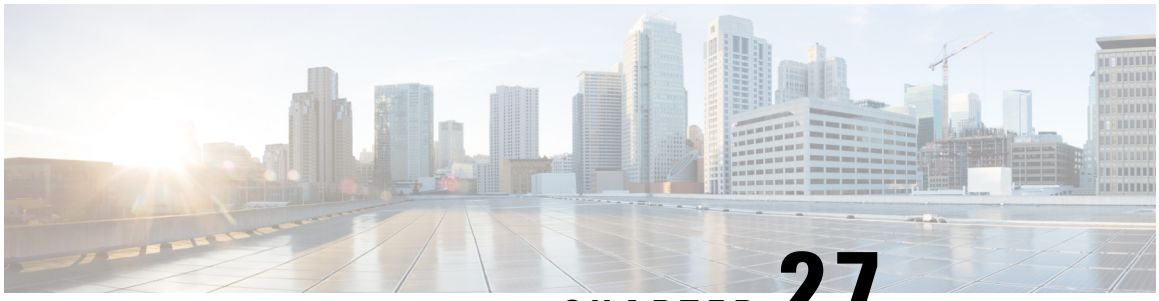
**Step 7** In the **Email Report To** field, enter an email ID or mailer ID if you want the report in an email.

You must configure SMTP settings in **Settings > Server Settings > SMTP** tab. If the Data service IP address is in private subnet, the static management route for SMTP server must be added in Cisco Nexus Dashboard cluster configuration.

**Step 8** In the **Select fabric(s)** table, select the MSD fabric on which the report job should be run.

**Step 9** Click **Save**.

The report job will be executed at the configured interval.



CHAPTER 27

# Configuring ToR Switches and Deploying Networks in External Fabrics

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


- Overview, on page 607
- Supported Topologies for ToR Switches, on page 607
- Configuring ToR Switches, on page 613
- Deploying Networks on ToR Switches, on page 615


## Overview

NDFC supports the Top-of-Rack (ToR) switches. You can add the Layer 2 ToR switches in an external fabric, and they can be connected to the Leaf switches in the Easy Fabric. Typically, the Leaf and ToR devices are connected with back-to-back vPC connection. For more information, see Supported Topologies for ToR Switches.

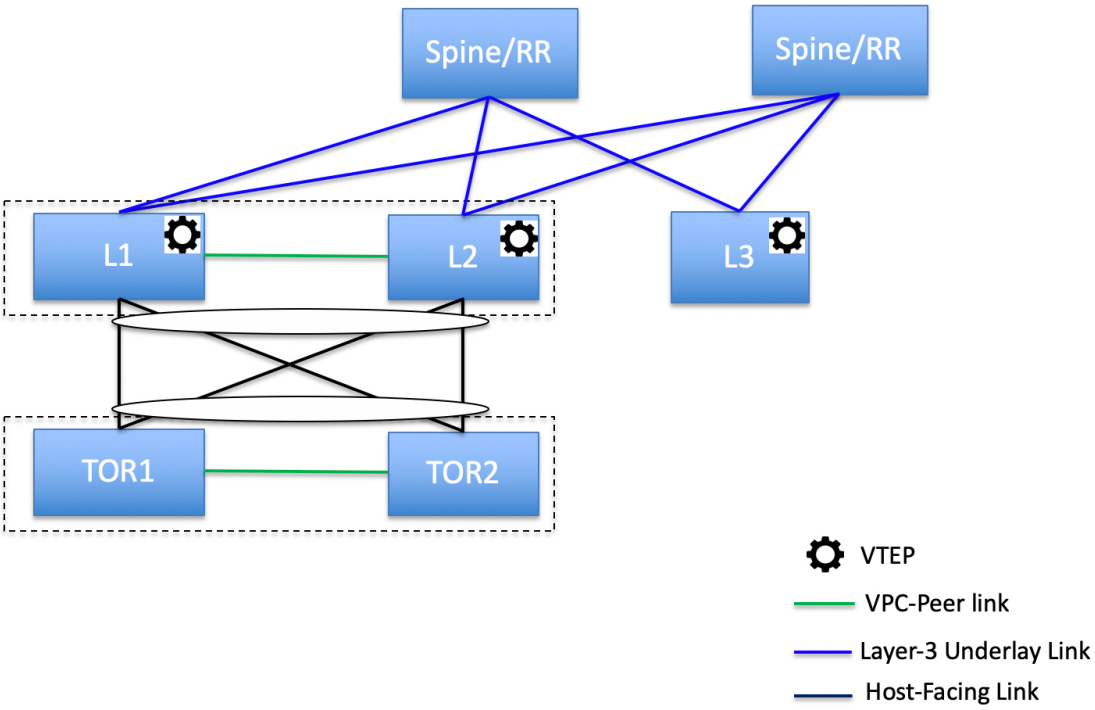
## Supported Topologies for ToR Switches

The following topologies with ToR switches are supported in NDFC:

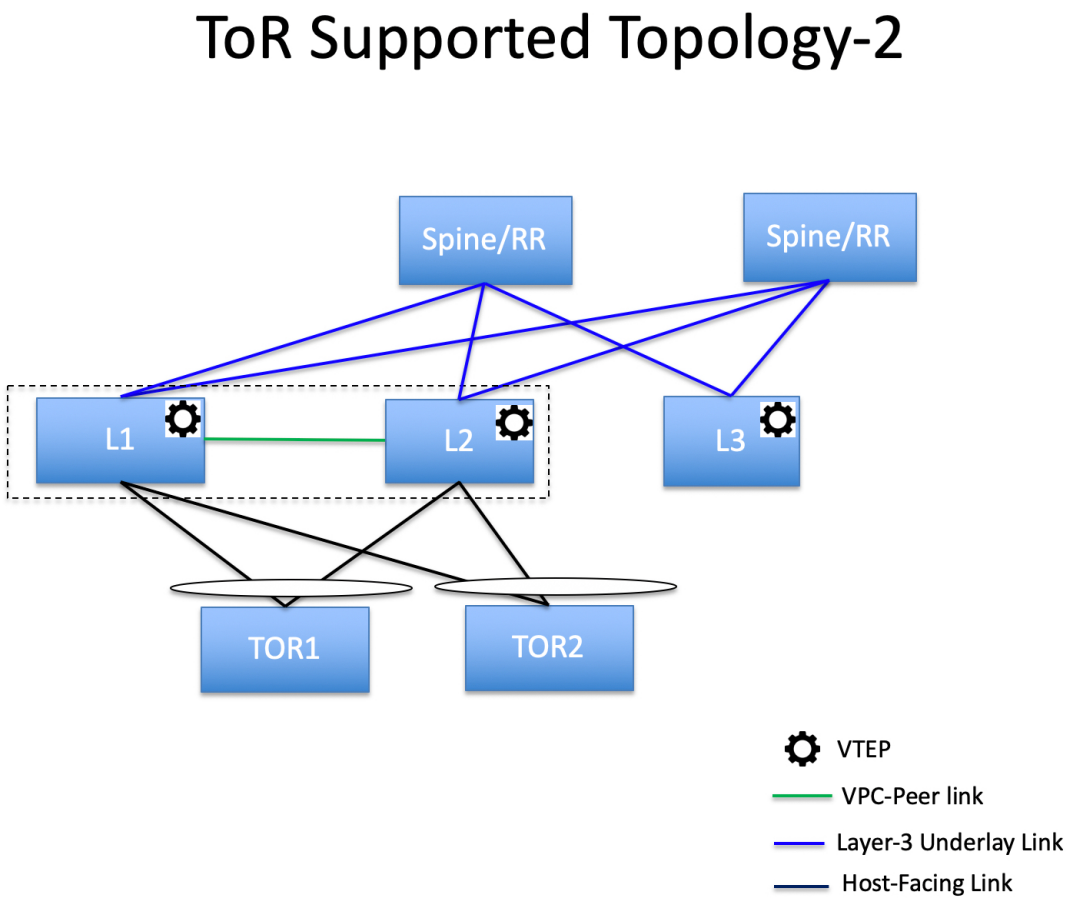
**Note** Cisco Nexus 7000 Series Switches do not support the **ToR** switch role in Cisco NDFC.

- ToR switches with back to back vPC connection to the leaf switches.

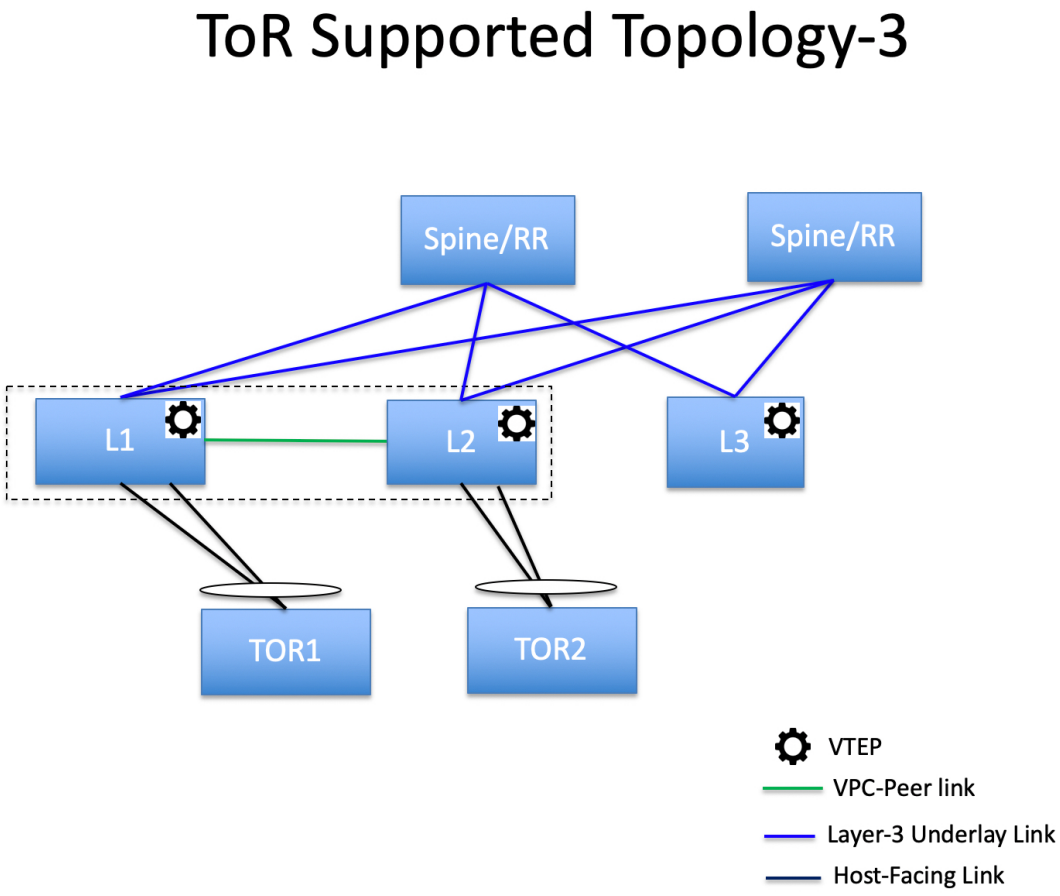
# ToR Supported Topology-1



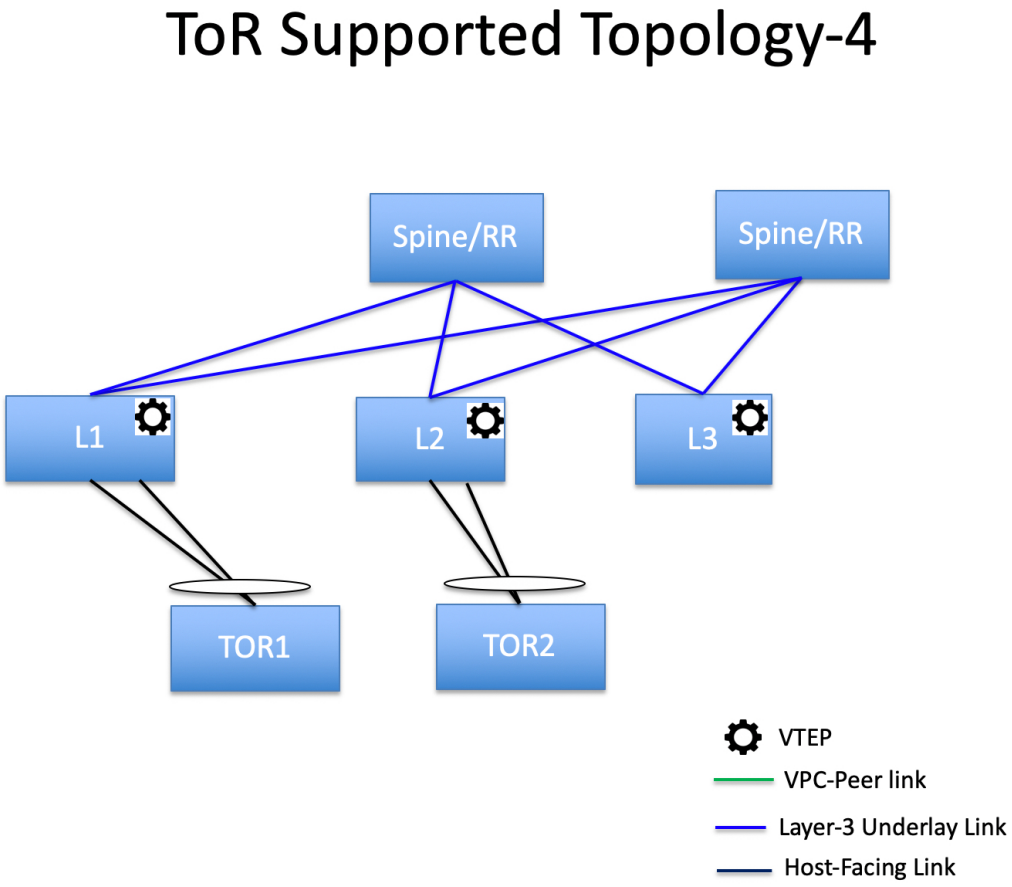

- ToR switches with port channels connected to both the leaf switches. The L1 and L2 switches are connected as a vPC pair.



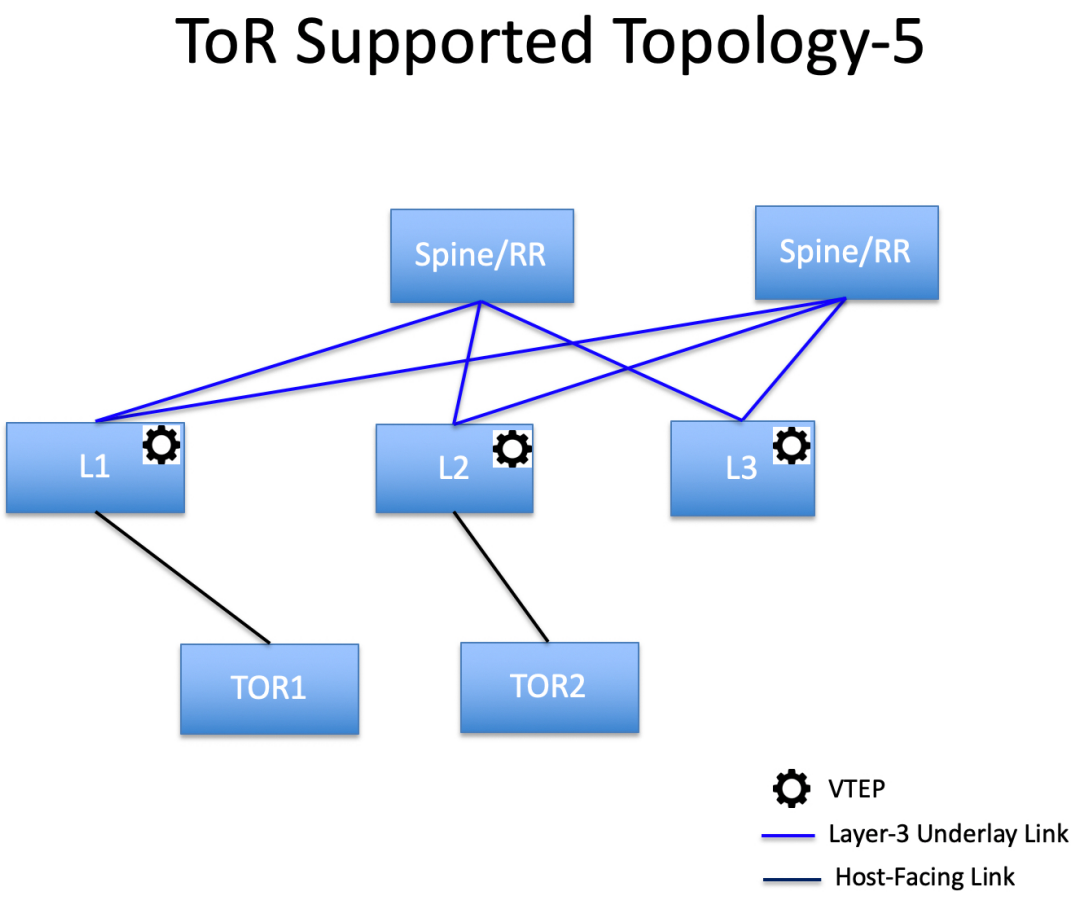

- ToR switches with port channels directly connected to the leaf switches. The L1 and L2 switches are connected as a vPC pair.



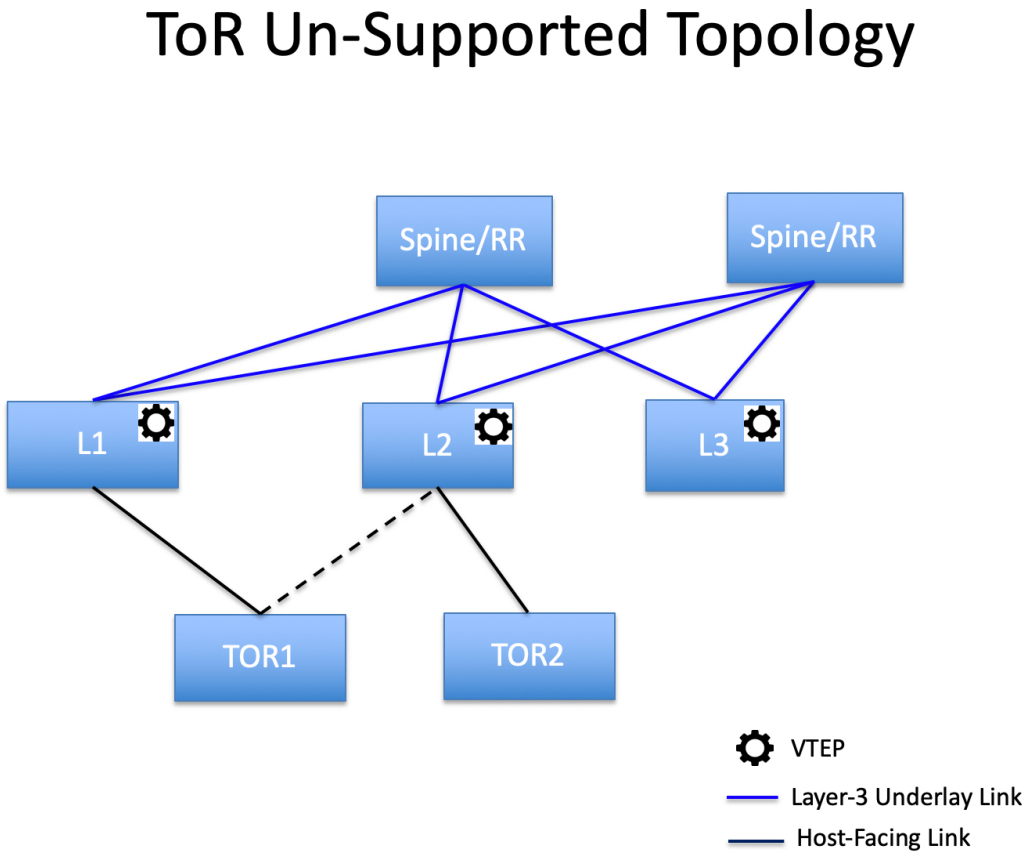

- ToR switches with port channels directly connected to the leaf switches. vPC pairs are not configured for the leaf or ToR switches.




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




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




# Configuring ToR Switches

Before you begin, make sure you have an Easy Fabric or create and deploy a new fabric. For more information, see Creating a VXLAN EVPN Fabric Using the Easy_Fabric Template, on page 49.

**Note** NDFC supports the trunk_host policies for the ToR switches. Make sure ToR has vPC, port channel or trunk host policy attached on the interfaces connected to Leaf. These policies are used to connect the ToR switches in the external fabric to the Leaf switches in the Easy Fabric.

**Procedure**

**Step 1** Create an external fabric and add two ToR switches. For more information, see Creating an External Fabric, on page 112.

The number of ToR switches can be more than two. This procedure shows how to configure ToR switches as shown in the ToR Topology-1, where ToR switches are connected using vPC. The following are the different scenarios for connecting the ToR switches:

- If vPC is not configured on the ToR switches, then vPC policy need to be applied on ToR facing interfaces if uplinks of these ToR switches are connected to vPC leaf switches.
- If ToR switches are connected to leaf using port-channel, then port-channel policies need to be applied on the ToR interfaces connected to the leaf switches.
- If ToR switches are connected to leaf switches as standalone, the trunk policies need to be applied on the TOR interfaces.

**Note**

- While creating the external fabric, make sure that the **Fabric Monitor Mode** check box is not chosen.
- The two ToR switches must be connected and have same switch role.

After adding the ToR switches, make sure that the role for the ToR switches is chosen as ToR.

**Step 2** Select one of the ToR switch and click on **Actions** > **vPC Pairing**.

Choose the second ToR switch as a vPC Peer.

**Step 3** Under vPC Pair Template, enter all the relevant details for a vPC connection between both the ToR switches. For more information about fields and their descriptions, see Creating a vPC Setup, on page 130.

**Note**
Step 2, and 3 are required since this example shows the ToR configuration for Topology-1. For Topology 2, 3, 4, and 5, the steps 2 and 3 are not required.

**Step 4** On **Switch Overview** window, click **Actions** > **Recalculate and Deploy**.

**Step 5** After the configuration is completed in the **Config Deployment** window, click **Close**.

**Step 6** Create the MSD fabric.

While creating the MSD fabric, under the General tab, choose the **ToR Auto-deploy Flag** check box. This action enables automatic deployment of the networks and VRFs in the Easy Fabric to the ToR switches in the External Fabric when you click **Recalculate and Deploy** in the MSD fabric. For more information, see Deploying Networks on ToR Switches, on page 615.

For information about the remaining tabs and fields, see *Creating an MSD Fabric*.

**Step 7** Open the MSD fabric. Navigate to **Child Fabrics** and click on **Actions** to move fabric into MSD. Select the Easy Fabric where ToRs are connected and click **Add**.

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

**Step 8** Open the Easy Fabric containing the leaf switches.

**Step 9** You need to create back-to-back vPC between the leaf and ToR switches.

**Step 10** Navigate to **LAN** > **Interfaces** > **Actions** > **Interface**.

Choose vPC and enter all the relevant details and click **Save**.

For more information about the fields in this window, see Adding Interfaces, on page 323.

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

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

# Deploying Networks on ToR Switches

To deploy networks on ToR switches in the external fabrics, you need to deploy them on the switches in the Easy Fabric through MSD. These switches should be connected to the ToR switches.

### Procedure

**Step 1** Choose **LAN** > **Fabrics**, double-click on the Easy Fabric.

**Step 2** In the **Networks** window, select the networks that you want to deploy or create a new network. For information about creating a network, see Creating Network for Standalone Fabrics, on page 214.

**Step 3** Select the **Network** from the **Network Attachment** window. Click on **Actions and Edit**. Attach the network and select the appropriate interface/port-channels and then click on **Save**. These port channels connect the leaf switches to the ToR switches. The networks will be deployed on these port channels.

**Step 4** On **Fabric Overview** window, click **Actions** > **Recalculate and Deploy**.

Now the VLANs are deployed on the leaf switches.

**Step 5** Navigate to MSD fabric.

**Step 6** On **Fabric Overview** window, click **Actions** > **Recalculate and Deploy**.

The networks created and deployed on the leaf switches in the Easy Fabric are also deployed on the ToR switches in the external fabric. This step allows the same VLANs to be configured on the ToR switches that are deployed on the leaf switches in the Step 4.

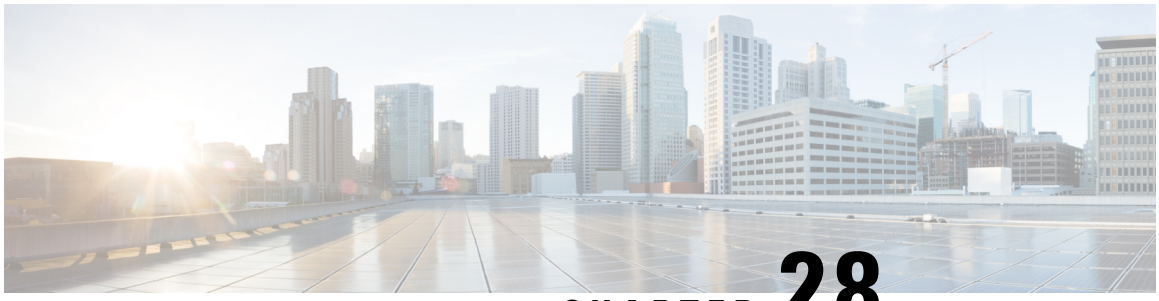
**Note**
If VLANs are created on the ToR switches manually using the freeform configs, they are not modified.



# PART VIII

# External/WAN Layer 3 Connectivity for VXLAN BGP EVPN Fabrics


- MPLS SR and LDP Handoff, on page 619
- VRF Lite, on page 629



CHAPTER 28

# MPLS SR and LDP Handoff

This chapter describes how to configure the MPLS handoff features.


- Overview of VXLAN EVPN to SR-MPLS and MPLS LDP Interconnection, on page 619
- VXLAN MPLS Topology, on page 621
- Configuration Tasks for VXLAN MPLS Handoff , on page 623
- Editing Fabric Settings for MPLS Handoff, on page 623
- Creating an Underlay Inter-Fabric Connection , on page 624
- Creating an Overlay Inter-Fabric Connection, on page 625
- Deploying VRFs, on page 626
- Changing the Routing Protocol and MPLS Settings, on page 627


## Overview of VXLAN EVPN to SR-MPLS and MPLS LDP Interconnection

Nexus Dashboard Fabric Controller (NDFC) supports the following handoff features:

- VXLAN to SR-MPLS
- VXLAN to MPLS LDP

These features are provided on the border devices, that is, border leaf, border spine, and border super spine in the VXLAN fabric using the **Data Center VXLAN EVPN** template. Note that the devices should be running Cisco NX-OS Release 9.3(1) or later. These DCI handoff approaches are the one box DCI solution where no extra Provider Edge (PE) device is needed in the external fabric.

**Note** If the switch is running a Cisco NX-OS Release 7.0(3)I7(X), enabling the MPLS handoff feature causes the switch to remove the NVE related config-profile CLIs when the switch is reloaded.

In the NDFC DCI MPLS handoff feature, the underlay routing protocol to connect a border device to an external fabric is ISIS or OSPF, and the overlay protocol is eBGP. The N-S traffic between the VXLAN fabric and external fabric running SR-MPLS or MPLS LDP is supported. Though, you can use NDFC for connecting two Data Center VXLAN fabrics via SR-MPLS or MPLS LDP.

### Supported Platforms and Configurations

The following table provides information about the supported platforms:

| Feature | Supported Platforms |
|---|---|
| VXLAN to SR-MPLS | Cisco Nexus 9300-FX2/FX3/GX, N9K-X96136YC-R, and Cisco Nexus 3600 R-Series switches |
| VXLAN to MPLS LDP | N9K-X96136YC-R and Cisco Nexus 3600 R-series switches |

The following features aren't supported as they aren't supported on a switch:

- Coexisting of MPLS LDP and SR-MPLS interconnections
- vPC

The VXLAN to SR-MPLS handoff feature comprises the following configurations:

- Base SR-MPLS feature configuration.
- Underlay configuration between the DCI handoff device and the device in the external fabric for the underlay connectivity. NDFC supports ISIS or OSPF as the routing protocol for the underlay connectivity.
- Overlay configuration between a DCI handoff device and a core or edge router in the external fabric, or another border device in another fabric. The connectivity is established through eBGP.
- VRF profile

The VXLAN to MPLS LDP handoff feature comprises the following configurations:

- Base MPLS LDP feature configuration.
- Underlay configuration between the DCI handoff device and the device in the external fabric for the underlay connectivity. NDFC supports ISIS or OSPF as the routing protocol for the underlay connectivity.
- Overlay configuration between a DCI handoff device and a core or edge router in the external fabric, or another border device in another fabric. The connectivity is established through eBGP.
- VRF profile

### Inter-Fabric Connections for MPLS Handoff

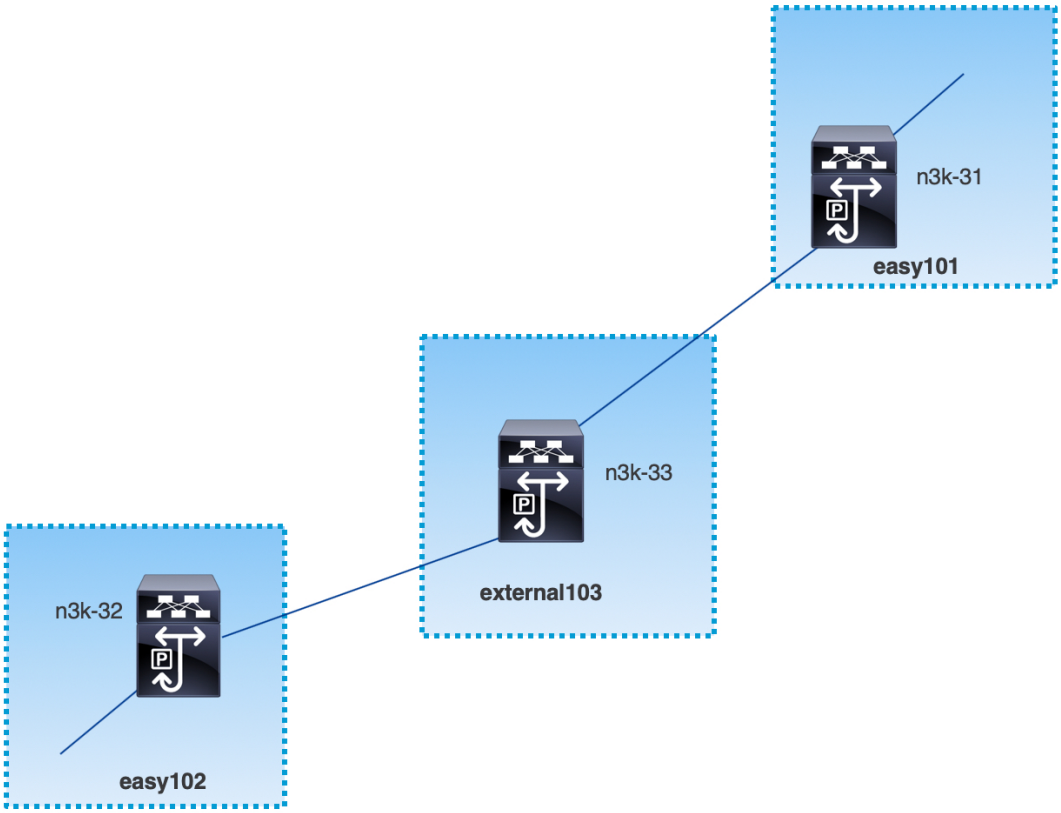
The following two inter-fabric connection links are introduced:

- **VXLAN_MPLS_UNDERLAY** for underlay configuration: This link corresponds to each physical link or Layer 3 port channel between the border and the external device (or a P router in MPLS or SR-MPLS). A border device can have multiple inter-fabric connection links as there could be multiple links connected to one or more external devices.
- **VXLAN_MPLS_OVERLAY** for eBGP overlay configuration: This link corresponds to the virtual link between a DCI handoff device and a core or edge router in the external fabric, or another border device in another fabric. This inter-fabric connection link can only be created on border devices which meet the image and platform requirement. A border device can have multiple of this type of IFC link as it could communicate to multiple core or edge routers.

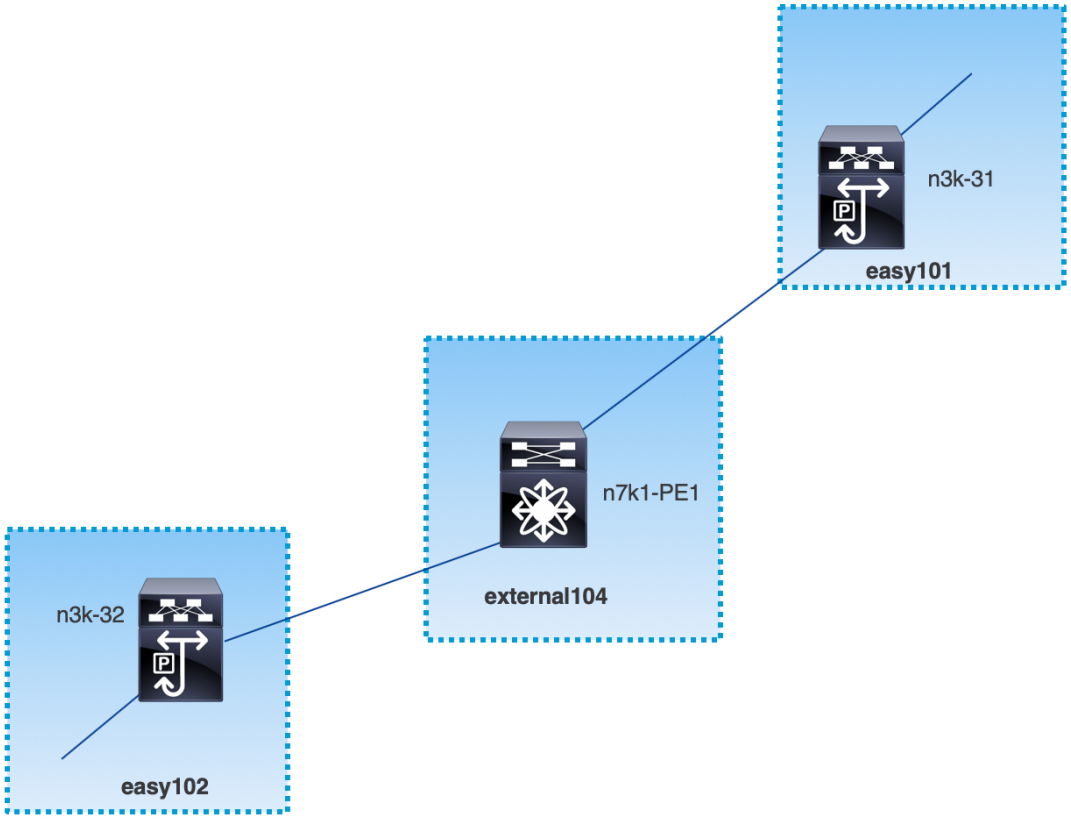
These inter-fabric connections can be manually created by using the NDFC Web UI or REST API. Note that the automatic creation of these inter-fabric connections isn't supported.

# VXLAN MPLS Topology

**MPLS-SR Topology**




**MPLS-LDP Topology**




This topology shows only the border devices in the Data Center VXLAN EVPN and the core or edge router in the External Connectivity Network fabric.

- The fabrics that are using the **Data Center VXLAN EVPN** template are:
  - **easy101**
  - **easy102**
- The fabrics that are using the **External Connectivity Network** template are:
  - **external103**
  - **external104**
- The external fabric **external103** is running the MPLS SR protocol.
- The external fabric **external104** is running the MPLS LDP protocol.
- **n3k-31** and **n3k-32** are border devices performing VXLAN to MPLS handoff.
- **n7k-PE1** only supports MPLS LDP.
- **n3k-33** supports SR-MPLS.

# Configuration Tasks for VXLAN MPLS Handoff

The following tasks are involved in configuring the MPLS handoff features:

1. Editing the fabric settings to enable MPLS handoff.
2. Creating an underlay inter-fabric connection link between the fabrics.

   Specify whether you're using MPLS SR or LDP in the inter-fabric connection link settings.
3. Creating an overlay inter-fabric connection link between the fabrics.
4. Deploying a VRF for VXLAN to MPLS interconnection.

# Editing Fabric Settings for MPLS Handoff

This section shows how to edit the fabric settings for the easy fabric and the external fabric to enable the MPLS handoff feature.

## Editing Easy Fabric Settings

**Procedure**

**Step 1** Choose **LAN** > **Fabrics**. Choose a appropriate fabric.

**Step 2** From **Actions** drop-down list, choose **Edit Fabric** to edit the fabric settings.

**Step 3** Click the **Advanced** tab.

**Enable MPLS Handoff**: Check the check box to enable the MPLS Handoff feature.

**Note**: For the brownfield import, choose the **Enable MPLS Handoff** feature. Most of the IFC configuration will be captured in **switch_freeform**.

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

**Step 4** Click the **Resources** tab.

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

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

**Step 5** Click **Save** to configure the MPLS feature on each border device in the fabric.

**Step 6** From **Actions** drop-down list, choose **Recalculate and Deploy**.

For more information about remaining fields, see Creating a New VXLAN BGP EVPN Fabric.

# Editing External Fabric Settings

**Procedure**

**Step 1** Choose **LAN** > **Fabrics**. Choose a appropriate fabric.

**Step 2** From **Actions** drop-down list, choose **Edit Fabric** to edit the fabric settings.

**Step 3** (Optional) Under the **General Parameters** tab, uncheck the **Fabric Monitor Mode** check box.

**Step 4** Click the **Advanced** tab.

**Enable MPLS Handoff**: Check the check box to enable the MPLS Handoff feature.

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

**Step 5** Click the **Resources** tab.

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

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

**Step 6** Click **Save** to configure the MPLS feature on each edge or core router in the fabric.

**Step 7** From **Actions** drop-down list, choose **Recalculate and Deploy**.

For more information about remaining fields, see Creating an External Fabric.

# Creating an Underlay Inter-Fabric Connection

This procedure shows how to create an underlay inter-fabric connection link.

**Procedure**

**Step 1** Choose **LAN** > **Fabrics**.

**Step 2** Choose a VXLAN fabric from which you want to create an underlay inter-fabric connection to MPLS.

**Step 3** On the **Fabric Overview** window, click the **Links** tab.

**Step 4** Check the existing links that are already discovered for the fabric.

In this example, the link from **easy101** to **external103** is already discovered.

**Step 5** Select the existing discovered link and click on **Actions** > **Edit**.

If a link isn't discovered, click on **Actions** > **Create** and provide all the details for adding an inter-fabric link.

**Step 6** In the **Link Management - Edit Link** window, provide all the required information.

**Link Type**: Choose inter-fabric.

**Link Sub-Type**: Choose VXLAN_MPLS_Underlay from the drop-down list.

**Link Template**: Choose ext_vxlan_mpls_underlay_setup from the drop-down list.

In the **General Parameters** tab, provide all the details.

**IP Address/Mask**: Specify the IP address with mask for the source interface.

**Neighbor IP**: Specify the IP address of destination interface.

**MPLS Fabric**: Specify whether the external fabric is running SR or LDP.

**Note**
MPLS SR and LDP can't coexist on a single device.

**Source SR Index**: Specify a unique SID index for the source border. This field is disabled if you choose **LDP** in the **MPLS Fabric** field.

**Destination SR Index**: Specify a unique SID index for the destination border. This field is disabled if you choose LDP for the **MPLS Fabric** field.

**SR Global Block Range**: Specify the SR global block range. You need to have the same global block range across the fabrics. The default range is from 16000 to 23999. This field is disabled if you choose LDP for the **MPLS Fabric** field.

**DCI Routing Protocol**: Specify the routing protocol used on the DCI MPLS underlay link. You can choose either **is-is** or **ospf**.

**OSPF Area ID**: Specify the OSPF area ID if you choose OSPF as the routing protocol.

**DCI Routing Tag**: Specify the DCI routing tag used for the DCI routing protocol.

**Step 7** Click **Save**.

**Step 8** On the **Fabric Overview** window, click on **Actions** > **Recalculate & Deploy**.

**Step 9** In the **Deploy Configuration** window, click **Deploy Config**.

**Step 10** Navigate to the destination fabric from the **LAN Fabrics** window and perform a **Recalculate & Deploy**, that is, perform steps 9 and 10.

# Creating an Overlay Inter-Fabric Connection

This procedure shows how to create an overlay inter-fabric connection after the underlay inter-fabric connection is created. The overlay inter-fabric connection is the same for MPLS SR and LDP because the overlay connection uses eBGP.

**Procedure**

**Step 1** On the **Links** tab, click on **Actions** > **Create**.

**Step 2** In the **Link Management - Create Link** window, provide all the details.

**Link Type**: Choose **Inter-Fabric**.

**Link-Sub Type**: Choose **VXLAN_MPLS_OVERLAY** from the drop-down list.

**Link Template**: Choose **ext_vxlan_mpls_overlay_setup** from the drop-down list.

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

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

**Source Device** and **Source Interface**: Choose the source device and the MPLS loopback interface. The IP address of the loopback interface will be used for overlay eBGP peering.

**Destination Device** and **Destination Interface**: Choose the destination device and a loopback interface that connects to the source device.

In the **General Parameters** tab, provide all the details.

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

**BGP Neighbor IP**: Fill up this field with the IP address of the loopback interface at the destination device for eBGP peering.

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

**Step 3** Click **Save**.

**Step 4** On the **Fabric Overview** window, click on **Actions** > **Recalculate & Deploy**.

**Step 5** In the **Deploy Configuration** window, click **Deploy Config**.

**Step 6** Navigate to the destination fabric from the **LAN Fabrics** window and perform a **Recalculate & Deploy**, that is, perform steps 4 and 5.

**Note**
If there is only one MPLS overlay IFC link on the switch, you can remove it only when there's no VRF attached to either end of the MPLS overlay link.

# Deploying VRFs

This procedure shows how to deploy VRFs for VXLAN to MPLS interconnection.



**Note** When you use the 4 byte ASN and auto route target is configured, the route target that is automatically generated is 23456:VNI. If two different VRFs in two different fabrics have the same VNI value, the route-target of the two VRFs would be the same due to auto route target and the value 23456 is always constant. For two fabrics connected via VXLAN MPLS handoff, this could result in unintended route exchange. Therefore, for security reasons, if you want to disable auto route target, you can disable it by customizing the network template and network extension template.

### Procedure

**Step 1** Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **VRFs**.

**Step 2** In the **VRFs** tab, click on **Actions** > **Create** to create a VRF. For more information, see Creating VRFs for the Standalone Fabric.

**Step 3** Select the newly added VRF and click **Continue**.

**Step 4** In the **VRF Deployment** window, you can see the topology of the fabric. Select a border device to attach a VRF to the border device where the MPLS LDP IFC link is created.

In this example, **n3k-31** is the border device in the **easy101** fabric.

**Step 5** In the **VRF Extension Attachment** window, select the VRF and click the **Freeform config** button under the CLI Freeform column.

**Step 6** Add the following freeform config manually to the VRF:

```
vrf context $$VRF_NAME$$
  address-family ipv4 unicast
    route-target import $$REMOTE_PE_RT$$
  address-family ipv6 unicast
    route-target import $$REMOTE_PE_RT$$
```

In the freeform config, *REMOTE_PE_RT* refers to the neighbor's BGP ASN and VNI number in the **ASN:VNI** format if the neighbor is a border device in Easy Fabric managed by NDFC.

**Step 7** Click **Save Config**.

**Step 8** (Optional) Enter the Loopback Id and Loopback IPv4 Address and IPv6 address for the border device.

**Step 9** Click **Save**.

**Step 10** (Optional) Click the **Preview** icon in the **VRF Deployment** window to preview the configuration that will be deployed.

**Step 11** Click **Deploy**.

Perform the same task from Step 3 to Step 11 in the destination fabric if the neighbor is a border device in Easy Fabric managed by NDFC.

# Changing the Routing Protocol and MPLS Settings

This procedure shows how to change the routing protocol of a device from using IS-IS to OSPF, or from using MPLS SR to LDP for underlay IFC.



**Note** MPLS SR and LDP cannot co-exist on a device, and using both IS-IS and OSPF for MPLS handoff on the same device is not supported.

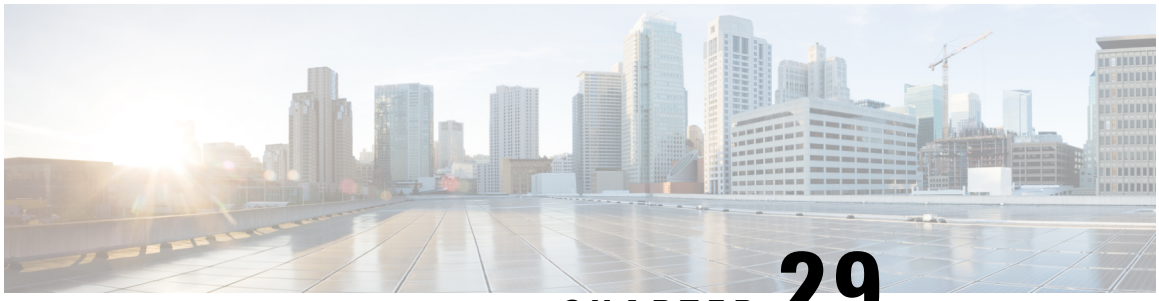
### Procedure

**Step 1** Remove all the MPLS underlay and overlay IFCs from the device that needs the change of DCI routing protocol or MPLS fabric.

**Step 2** Click **Recalculate & Deploy** for each fabric that is involved in the removal of the IFCs.

This step deletes all global MPLS SR/LDP configurations and the MPLS loopback interface that was previously created.

**Step 3** Create a new IFC using the preferred DCI routing protocol and MPLS settings. For more information, see Creating an Underlay Inter-Fabric Connection , on page 624.



CHAPTER 29

# VRF Lite


- VRF Lite, on page 629
- Prerequisites and Guidelines, on page 630
- Sample Scenarios, on page 631
- Automatic VRF Lite (IFC) Configuration, on page 631
- VRF Lite Between Cisco Nexus 9000 Based Border and Cisco Nexus 9000 Based Edge Router, on page 633
- VRF Lite Between Cisco Nexus 9000 Based Border and Non-Cisco Device, on page 638
- VRF Lite Between Cisco Nexus 9000 Based Border and Non-Nexus Device, on page 642
- Appendix , on page 643


## VRF Lite

External connectivity from data centers is a prime requirement where workloads that are part of a data center fabric can communicate with an outside fabric over WAN or Backbone services. To enable Layer-3 for north-south traffic flow, use virtual routing and forwarding instances (VRF) Lite peering between data center border devices and the external fabric edge routers.

A VXLAN (Virtual extensible Local Area Network) EVPN (Ethernet Virtual Private Network) based data center fabrics provide connectivity by distributing IP-MAC reachability information among various devices within the fabric. The VRF Lite feature is used for connecting the fabric to an external Layer 3 domain. This can be a border router or a Border Gateway router.

You can enable VRF-Lite on the following devices:

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

# Prerequisites and Guidelines

- VRF Lite requires Cisco Nexus 9000 Series Cisco Nexus Operating System (NX-OS) Release 7.0(3)I6(2) or later.
- Familiarity with VXLAN BGP EVPN data center fabric architecture and VXLAN Overlays provisioning through NDFC.
- Fully configured VXLAN BGP EVPN fabrics including underlay and overlay configurations for the various leafs and spine devices, external fabric configuration through NDFC, and relevant external fabric device configuration (edge routers, for example).
    - You can configure VXLAN BGP EVPN fabric (and its connectivity to an external Layer 3 domain for north-south traffic flow) manually or using NDFC.

      This document explains the process to connect the fabric to an edge router (outside the fabric, toward the external fabric) through NDFC. So, you must know how to configure and deploy VXLAN BGP EVPN and external fabrics through NDFC.
    - VRF Lite can be enabled on physical Ethernet interface or Layer 3 port-channel. Subinterface over physical interface or Layer 3 port-channel interface that is created in NDFC at the VRF extension moment for each VRF lite link that the VRF is extended over.
- To delete a VRF Lite IFC, remove all VRF extensions that are enabled on the IFC. Else, an error message is reported. After you remove the VRF Lite attachments, recalculate and deploy the fabric to remove any pending Layer-3 extension configurations. It removes the per-VRF subinterface and per-VRF External Border Gateway Protocol configuration on the devices.
- When you create a VXLAN VRF, ensure that you check the following 3 fields:
    - **Advertise Host Routes** – By default, over the VRF Lite peering session, only nonhost (/32 or /128) prefixes are advertised. If host routes (/32 or /128) must be enabled and advertised from the border device to the edge/WAN router, check the **Advertise Host Routes** check box. Route-map does outbound filtering. By default, this check box is disabled.
    - **Advertise Default Route** – This field controls whether a network statement 0/0 will be enabled under the VRF. This in turn advertises 0/0 route in BGP. By default, this field is enabled. When you choose this check box, this ensures that a 0/0 route is advertised inside the fabric over EVPN Route-type 5 to the leafs, there by providing a default route out of the Leafs toward the border devices.
    - **Config Static 0/0 Route** – By default, this check box is checked. This field controls whether a static 0/0 route to the edge/WAN router, must be configured under the VRF on the border device. By default, this field is enabled. If WAN/edge routers are advertising a default route over the VRF Lite peering to the border device in the fabric, then this field must be disabled. In addition, the **Advertise Default Route** field must be disabled. The 0/0 route that is advertised over External Border Gateway Protocol sends over EVPN to the leafs without requirement of more configuration. The clean iBGP EVPN separation inside the fabric with eBGP for external out-of-fabric peering provides provides for this desired behavior.

# Sample Scenarios

The following sections explain different use-cases for configuring VRF Lite:

- Automatic VRF Lite (IFC) Configuration
- VRF Lite between Cisco Nexus 9000 based Border and Cisco Nexus 9000 based Edge Router
- VRF Lite between Cisco Nexus 9000 based Border and Non-Cisco device
- VRF Lite between Cisco Nexus 9000 based Border and Non-Nexus device

  This is a typical use-case of Cisco ASR 9000 based Edge Router in Managed mode

# Automatic VRF Lite (IFC) Configuration

**Guidelines**

- Auto IFC is supported on Cisco Nexus devices only.
- You can configure Cisco ASR 1000 Series routers and Cisco Catalyst 9000 Series switches as edge routers. To configure, set up a VRF Lite IFC, and connect it as a border device with easy fabric.
- You can configure Cisco ASR 9000 Series routers as edge routers in managed mode.
- If the device in the External fabric is non-Nexus, you must create IFC manually.
- Ensure that no user policy is enabled on the interface that connects to the edge router. If a policy exists, then the interface will not be configured.
- Autoconfiguration is supported for the following cases:
  - **Border** role in the VXLAN fabric and **Edge Router** role in the connected external fabric device
  - **Border Gateway** role in the VXLAN fabric and **Edge Router** role in the connected external fabric device
  - **Border** role to another **Border** role directly



**Note** Autoconfiguration is not provided between two Border Gateways (BGWs).

If VRF Lite is required between other roles, you must deploy it manually on the NDFC Web UI.

- To deploy configurations in the external fabric, you must uncheck the **Fabric Monitor Mode** check box in the external fabric settings. When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on the switches.

**Easy Fabric Settings**

The following are the 2 modes in which you can deploy VRF Lite. By default, VRF Lite deployment is set to Manual. You can change the settings based on your requirement.

- **Manual** - Use this option to deploy the VRF Lite IFCs manually between the source and the destination devices.
- **Back2Back&ToExternal** - Use this option to automatically configure VRF Lite IFCs between a border switch and the edge or core switches in external fabric or between back-to-back border switches in VXLAN EVPN fabric.



**Note** Though VRF Lite mode is set to **Manual** for NDFC resource handling, Data Center Interconnectivity (DCI) subnet is required.

The **Manual** mode is the default mode in fabric settings. To change the default mode to other mode, click **Edit** fabric settings. On **Resource** tab, modify **VRF Lite Deployment** field to the above mentioned auto configuration modes.

**Auto Deploy Peer** - This check box is applicable for VRF Lite deployment. When you check this check box, IFCs are automatically created for peer devices. You can check or uncheck this check box when the **VRF Lite Deployment** field is not set to **Manual**. The value you choose takes priority. This configuration only affects the new auto-created IFCs and does not affect the existing IFCs.

**Auto Deploy Default VRF** – When you select this check box, the **Auto Generate Configuration on default VRF** field is automatically enabled for auto-created VRF Lite IFCs. You can check or uncheck this checkbox when the VRF Lite Deployment field is not set to **Manual**. The **Auto Generate Configuration on default VRF** field when set, automatically configures the physical interface for the border device in the default VRF, and establishes an EBGP connection between the border device and the edge device or another border device in a different VXLAN EVPN fabric.

**Auto Deploy Default VRF for Peer** – When you select this checkbox, the **Auto Generate Configuration for NX-OS Peer on default VRF** field is automatically enabled for auto-created VRF Lite IFCs. You can check or uncheck this checkbox when the **VRF Lite Deployment** field is not set to **Manual**. The **Auto Generate Configuration for NX-OS Peer on default VRF** field when set, automatically configures the physical interface and the EBGP commands for the peer NX-OS switch.



**Note** You can access the **Auto Generate Configuration on default VRF** and **Auto Generate Configuration for NX-OS Peer on default VRF** fields for an IFC link by navigating to **Actions** > **Edit** > **VRF Lite**.

**Redistribute BGP Route-map Name** – Defines the route map for redistributing the BGP routes in default VRF.

**VRF Lite Subnet IP Range**: The IP address for VRF Lite IFC deployment is chosen from this range. The default value is 10.33.0.0/16. Ensure that each fabric has its own unique range and is distinct from any underlay range to avoid possible duplication. These addresses are reserved with the Resource Manager.

**VRF Lite Subnet Mask**: By default, it is set to /30, which is a best practice for point-to-point (P2P) links.

# VRF Lite Between Cisco Nexus 9000 Based Border and Cisco Nexus 9000 Based Edge Router

DC-Vxlan VXLAN EVPN Fabric is connected to WAN-Vxlan cloud. In the following topology, you can view WAN-Vxlan. The easy fabric has border leaf role and WAN-Vxlan cloud has a device with role edge router. NDFC shows physical and logical representation of the topology with CDP/LLDP Link discovery.

In this example, you can enable VRF Lite connections between DC-Vxlan border leaf and WAN-Vxlan edge router.

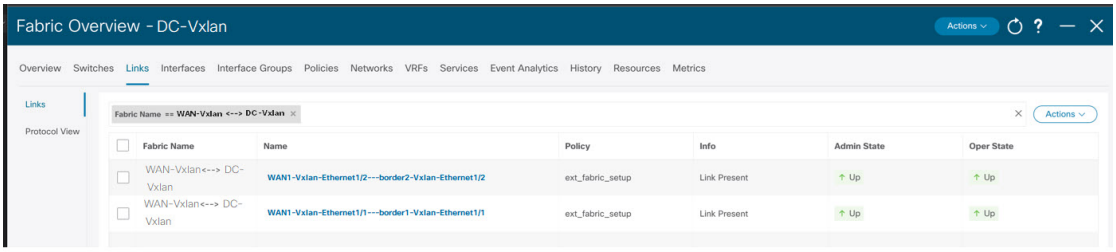
For VRF Lite configuration, you must enable External Border Gateway Protocol (EBGP) peering between the fabric's border interfaces and the edge router's interfaces, through point-to-point (P2P) connections.

The border physical interfaces are:

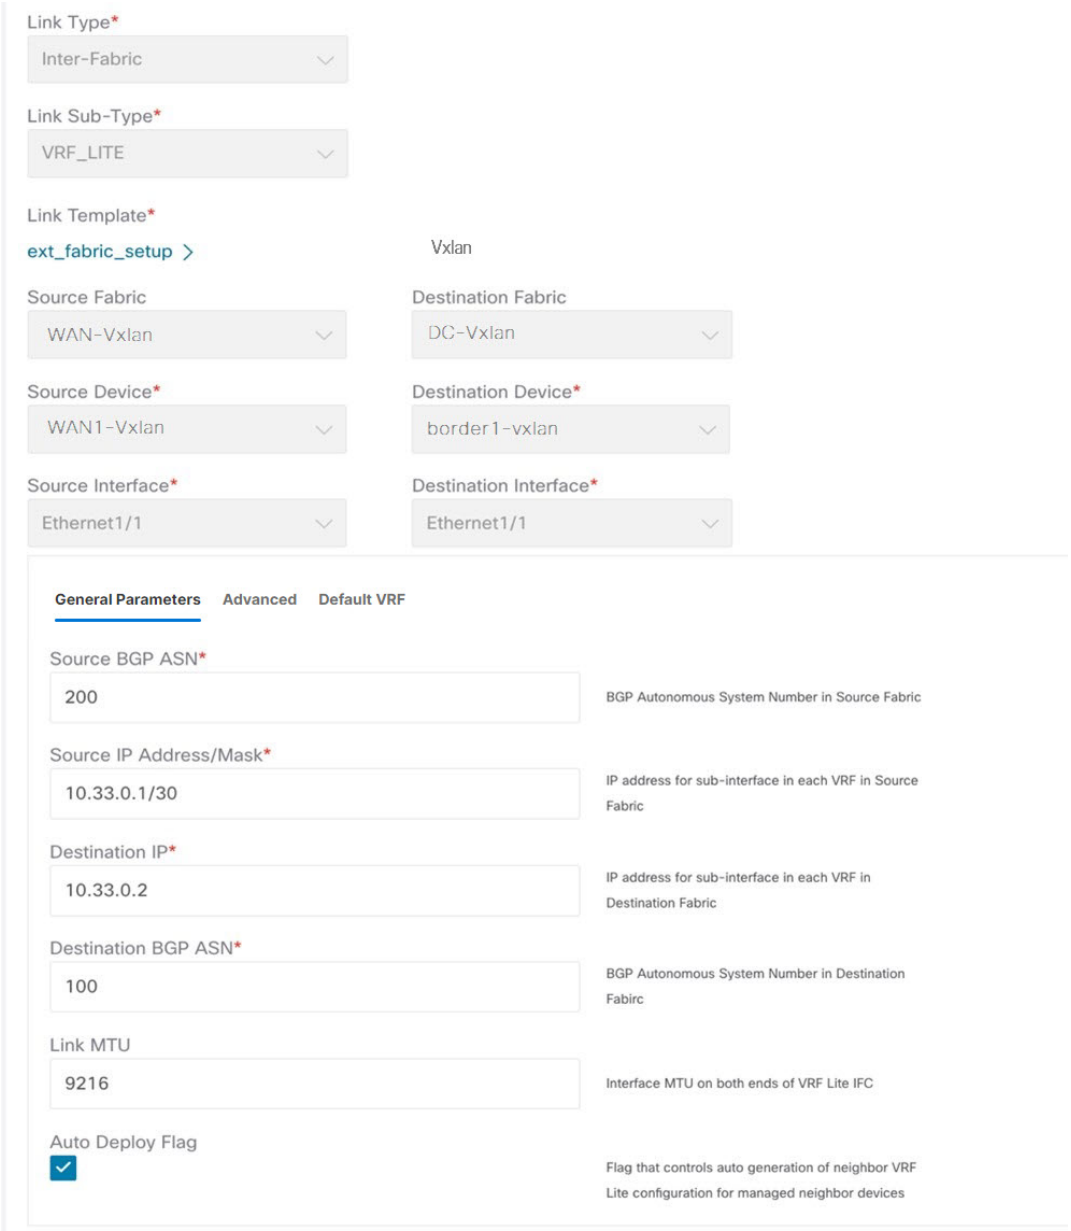
- **eth1/1** on **border1-Vxlan**, toward **eth1/1** on **WAN1-Vxlan**.
- **eth1/2** on **border2-Vxlan**, toward **eth1/2** on **WAN1-Vxlan**.

1. Verify the links between the border and the edge router. Choose **LAN** > **Fabrics**, double-click on **DC-Vxlan** fabric.

   On **Fabric Overview** window, click **Links** tab. You can view that the links that are detected by NDFC and **ext_fabric_setup** policy are assigned automatically.




2. Perform the following to verify the VRF Lite configuration:
   a. Select the fabric name and choose **Actions** > **Edit**.
   b. Click appropriate **Links**, choose **Actions** > **Edit**.




c. **Link Type** – Specifies the Interfabric link between two different fabrics within NDFC.

d. **Link Sub-Type** – Specifies the subtype of link. By default, the **VRF_LITE** option is displayed.

e. **Link Template** – Specifies the template for the link. The default template for a VRF Lite IFC is **ext_fabric_setup** is displayed. The template enables the source and destination interfaces as Layer 3 interfaces, figures the **no shutdown** command, and sets their MTU to 9216.

   The Source and Destination Fabric, Device, and Interfaces are autodetected and chosen by NDFC based on CDP/LLDP discovery.

f. On the **General Parameters** tabs, the fields in this tab are:

   - **Source BGP ASN** – BGP ASN of selected source fabric

- **Source IP Address/Mask** - NDFC auto allocated IP pool from Resource Manager Pool of VRF Lite subnet Pool for the **Ethernet1/1** subinterfaces, the source interface of the IFC. A subinterface is created for each VRF extended over this IFC, and a unique 802.1Q ID is assigned to it. The IP address/Mask entered here, along with the BGP Neighbor IP field (explained below) will be used as the default values for the subinterface that is created at VRF extension and can be overwritten.

  For example, an 802.1Q ID of 2 is associated with subinterface Eth 1/1.2 for VRF CORP traffic, and 802.1Q ID of 3 is associated with Eth 1/1.3 and VRF ENG, and so on.

  The IP prefix is reserved with the NDFC resource manager. Ensure that we use a unique IP address prefix for each IFC we create in the topology.

- **Destination IP** - NDFC auto allocated IP pool from resource manager pool of VRF Lite subnet pool. This is a BGP neighbor IP on the device.

  Interfabric traffic from different VRFs for an IFC 's the same source IP address (10.33.0.1/30) and destination IP address (10.33.0.2) as an example.

- **Destination BGP ASN** – BGP ASN of selected Destination fabric
- **Link MTU** – Default 9216
- **Auto Deploy Flag** – Default Auto selected based on fabric settings. This knob autoconfigures the neighbor VRF on neighboring managed device. For example, it will automatically create VRF on the edge router inside WAN-Vxlan External fabric.

**g.** The **Advanced** tab is added in the **Link Profile** section. The fields in this tab are:

- **Source Interface Description**
- **Destination Interface Description**
- **Source Interface Freeform Config**
- **Destination Interface Freeform Config**

**h.** Enter the following details in the **Default VRF** tab.

- **Auto Generate Configuration on default VRF**: Automatically configures the physical interface for the border device in default VRF. Establishes an EBGP connection between the border device and its peer in the VRF LITE IFC (edge router or another border device in a different VXLAN EVPN fabric).
- **Auto Generate Configuration for NX-OS Peer on default VRF**: Automatically configures the physical interface and the BGP commands for the NX-OS peer switch corresponding to default VRF.
- **Redistribute BGP Route-map Name**: Defines the route map used for redistributing the BGP routes in default VRF.
- **Default VRF BGP Neighbor Password**: Allows you to provide additional security by specifying a password for the BGP neighbor.
- **Default VRF BGP Password Key Encryption Type**: Specifies the encryption type for the password.

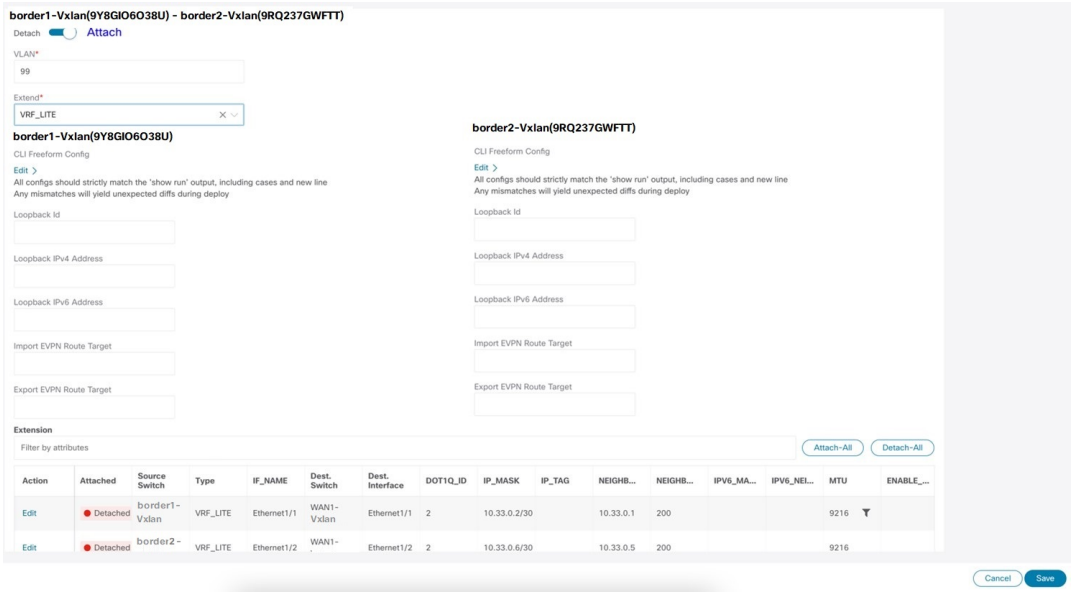
- **VRF Name in NX-OS Peer Switch**: Enables you to provide a name for the VRF in the peer switch. By default, the peer device is configured with default VRF, if left blank.
- **Enable DCI tracking**: Enables Data Center Interconnectivity (DCI) tracking on the physical interface, which allows you to use an existing VRF Lite link as the VXLAN EVPN multi-site underlay link. Note that you must enable this option only on a border gateway device.
- **Routing TAG**: Specifies the routing tag associated with the interface IP.

**Note** All the configuration changes made to the devices as a result of the above configurations are displayed in the **Pending Config** window.

i. Click **Save** to save the configuration.

3. To attach VRF and VRF Lite extensions on the border devices:
   a. Click **VRFs** > **VRF Attachments** tab.
   b. Choose **VRF Name**, click **Actions** > **Edit**.

      The **Edit** window appears.
   c. You can edit details in **Extension** field as mentioned below:




- Toggle the knob to **Attach**.
- In **Extend**, choose **VRF_LITE** from the drop-down list.
- On **Extension** card, choose one switch at a time, click **Edit**, enter details for **PEER_VRF_NAME**. This auto deploys the VRF on the neighboring device.

When you extend VRF Lite consecutive scenario, the VRF must be in the peer fabric and VRF name must be same. If the VRF is not in the peer fabric and if you try to extend VRF Lite, an error message is generated displaying the issue.

When you extend VRF Lite between an easy fabric and an external fabric, the VRF name can be same as name of source fabric, or default name, or another VRF name. Enter required VRF name in **PEER_VRF_NAME** field. The child PTIs for subinterface, VRF creation and BGP peering on external fabric have source values that are populated in it, hence the policies cannot be edited or deleted.

Follow above procedure for other links.

On **Edit** window, click **Attach-All**, to attach the required VRF Extension on the border devices, and then click **Save**.

4. To Recalculate and deploy configurations on VXLAN EVPN Easy Fabric:

   On **Fabric** windows double-click on appropriate fabric to navigate to **Fabric Overview** window. Click **Actions** > **Recalculate & Deploy**.

   Similarly, you can also perform operation, choose required **VRF Name** on **VRF attachments** tab, click **Actions** > **Deploy** to initiate VRF and VRF Lite configurations on the border devices.

5. To Recalculate and Deploy VXLAN EVPN Easy fabric:

   On **Fabric** window, click **Action** > **Recalculate and Deploy**.

   Similarly, you can choose the VRF attachments, edit, and click **Deploy**. It pushes VRF and VRF Lite configurations the border devices.

6. To recalculate and deploy configurations on external fabric, choose external fabric and follow the above procedure.

# VRF Lite Between Cisco Nexus 9000 Based Border and Non-Cisco Device

This example displays the procedure to enable VRF Lite connections between the DC-VXLAN border leaf device and a non-Cisco device in external fabric.

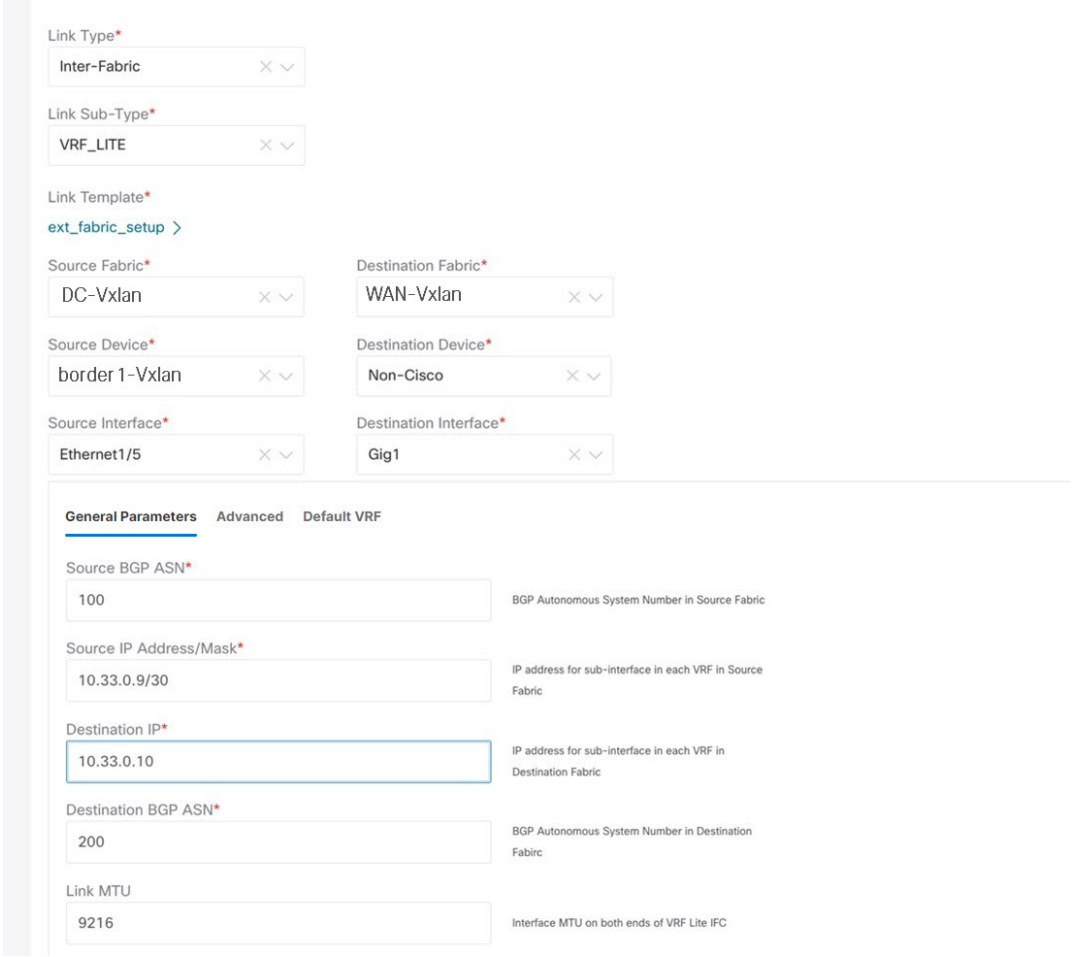
It is recommended to use meta definition of a device instead of importing devices in external fabric. This allows VRF Lite configurations to extend Cisco Nexus 9000 managed border devices in easy fabric. NDFC will not manage destination non-Cisco device. You must configure relevant VRF Lite configuration on the destination device.

1. To create new IFC links between border and edge router.
   a. On **Fabrics** window, double-click the fabric.

      The **Fabric Overview** window appears.

   b. Navigate to **Links** tab. On **Links** tab, click **Actions** > **Create a new link**.

      The **Create New link** window appears.




**c.** Enter the following required parameters in the window:

- **Link Type** – Select the Interfabric link. This is the IFC between two different fabrics within NDFC.
- **Link Sub-Type** - By default, the **VRF_LITE** option is displayed.
- **Link Template** – The default template for a VRF Lite IFC, **ext_fabric_setup**, is displayed. The template enables the source and destination interfaces as Layer 3 interfaces, configures the **no shutdown** command, and sets their MTU to 9216.
- **Source Fabric** – Select the Source Fabric. This is the Easy fabric where Cisco Nexus 9000 based border device resides.
- **Destination Fabric** – Select any External or Classic LAN fabric. It can be monitor mode as well.
- **Source Device** – Select the Source Device. This is the Cisco Nexus 9000 based border device.
- **Destination Device** – Now, you can create a “meta device definition”. Type any name and click create. For example, non-cisco.
- **Source Interface** – Select the interface on the border device where the non-cisco device is connected.

- **Destination Interface** – Now, you can create a “meta device interface”. Type any interface name and click create. For example, gig1, tengig1/10, eth1/1 are the valid interface names.

The **General Parameters** tab has the following fields:

- **Source BGP ASN** – BGP ASN of selected Source fabric.
- **Source IP Address/Mask -** Provide IP address and mask for the **Ethernet1/5** subinterfaces, the source interface of the IFC. Subinterface is created for each VRF extended over this IFC, and a unique 802.1Q ID is assigned to it. The IP address/Mask entered here, along with the BGP Neighbor IP field (explained below) used as the default values for the subinterface that is created at VRF extension and can be overwritten.

  For example, an 802.1Q ID of 2 is associated with subinterface Eth 1/5.2 for VRF CORP traffic, and 802.1Q ID of 3 is associated with Eth 1/5.3 and VRF ENG, and so on.

  The IP prefix is reserved with the NDFC resource manager. Ensure that we use a unique IP address prefix for each IFC we create in the topology.
- **Destination IP -** NDFC auto allocated IP pool from Resource Manager Pool of VRF Lite subnet Pool. It is a BGP neighbor IP on the device.

  Interfabric traffic from different VRFs for an IFCs the same source IP address (10.33.0.1/30) and destination IP address (10.33.0.2) as an example.
- **Destination BGP ASN** – BGP ASN of selected Destination fabric
- **Link MTU** – Default 9216
- **Auto Deploy Flag** – Not applicable as the destination device is Non-Nexus and Non-Cisco.

Enter the appropriate details in the **Advanced** tab. The following mentioned fields are in the tab:

- **Source Interface Description**
- **Destination Interface Description**
- **Source Interface Freeform Config**
- **Destination Interface Freeform Config**
- **Template for Configuration Generation on Peer**

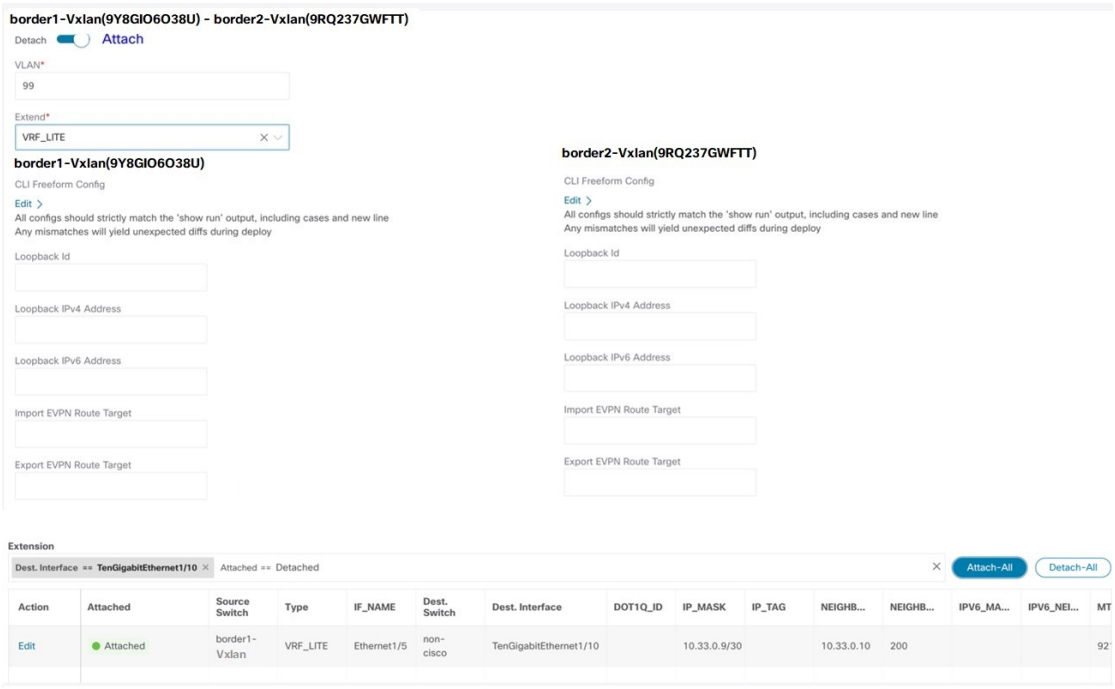
Enter the following details in the **Default VRF** tab.

- **Auto Generate Configuration on default VRF**: Automatically configures the physical interface for the border device in default VRF. Establishes an EBGP connection between the border device and its peer in the VRF Lite IFC (edge router or another border device in a different VXLAN EVPN fabric).
- **Auto Generate Configuration for NX-OS Peer on default VRF**: This is not applicable as the peer is a non NX-OS device.
- **Redistribute BGP Route-map Name**: Defines the route map used for redistributing the BGP routes in default VRF.
- **Default VRF BGP Neighbor Password**: Allows you to provide additional security by specifying a password for the BGP neighbor.

- **Default VRF BGP Password Key Encryption Type**: Specifies the encryption type for the password.
- **VRF Name in NX-OS Peer Switch**: Enables you to provide a name for the VRF in the peer switch. By default, the peer device is configured with default VRF, if left blank.
- **Enable DCI tracking**: Enables Data Center Interconnectivity (DCI) tracking on the physical interface, which allows you to use an existing VRF link as the VXLAN EVPN multi-site underlay link. Note that you must enable this option only on a border gateway device.
- **Routing TAG**: Specifies the routing tag associated with the interface IP.

**Note** All the configuration changes made to the devices as a result of the above configurations are displayed in the **Pending Config** window.

2. Click **Save** to create new link with parameters mentioned.
3. To attach VRF and VRF Lite extensions on the border devices, double-click on **DC-Vxlan** fabric. On **Fabric Overview** window, navigate to **VRFs** > **VRF Attachments** and edit the details as shown in the following image.




Click **Attach-all** to attach the required VRF Extension on the border devices and then click **Save**.

4. To recalculate and deploy configurations on VXLAN EVPN Easy fabric, click appropriate fabric on **Fabric** window.

On **Fabric Overview** window, click **Actions** > **Recalculate & Deploy**, or navigate to **VRF** > **VRF attachments**, choose the VRF attachments, edit, and then click **Deploy**. This initiates the VRF and VRF Lite configurations on the border devices.

# VRF Lite Between Cisco Nexus 9000 Based Border and Non-Nexus Device

In this example, you can enable VRF Lite connections between DC-Vxlan border leaf and a non-Nexus device in an external fabric.

Before Cisco NDFC Release 12.0.1a, ASR 9000 was supported for external fabric in monitor mode only. From Release 12.0.1a, ASR 9000 is supported in managed mode with an edge router role.

The following are the supported platforms:

- ASR 9000
- NCS 5500
- ASR 8000

Configuration compliance is enabled for IOS-XR switches in external fabric, similar to Cisco Nexus switches configured on external fabric. NDFC pushes configuration at the end of deployment

**Note** Ensure that the VXLAN BGP EVPN border device is active.

### Procedure

**Step 1** Navigate to **LAN** > **Fabrics** to create external fabric.

**Step 2** On **Create Fabric** window, enter appropriate ASN number, uncheck **monitor mode** check box, and then click **Save**.

**Step 3** Navigate to **Switches** window, click **Actions** > **Add switches**.

**Note**
Ensure that the IOS-XR device has the IP address reachability to NDFC with SNMP configurations for discovery.

To add non-Nexus devices to external fabrics, see Adding Non-Nexus Devices to External Fabrics, on page 121.

**Step 4** On **Add Switches** window, choose **Discover** check box, and **IOS-XR** from drop-down list for **Device Type** field.

**Step 5** After the router is discovered, you can view the switch name in the **Discovery Results** field.

**Step 6** Choose the discovered router and add to fabric. Ensure that the **Discovery Status** displays **OK** in the status column. Edge router role is supported.

After successful discovery, you can view the links between the devices in the **Links** tab.

**Step 7** To create VRF Lite IFC for external fabric with Cisco Nexus 9000 border leaf, choose the link and click **Actions** > **Edit**.

**Step 8** On **Edit Link** window, fill the required details for IFC creation. Few fields are auto-populated.

**Note**
For non-NX-OS device auto, deploy flag is not applicable.

**Step 9** To extend VRF Lite configurations on VXLAN border device, navigate to **VRF** > **VRF Attachment** tab, choose the VRF name, click **Actions** > **Edit** and then extend it as VRF Lite.

**Step 10** Deploy the configuration on VXLAN border device.

**Step 11** Navigate to the fabrics window, ensure that the external fabric has the router, click **Apply** to VRF Lite BGP policies.

**Step 12** Navigate to **Policies** tab, and add policies **ios_xr_base_bgp** and enter required details and click **Save**.

**Step 13** Add another policy **ios_xr_Ext_VRF_Lite_Jython** and enter required details and click **Save**

**Step 14** Deploy the configurations on the IOS-XR router.

# Appendix

### Nexus 9000 Border device configurations

Border-Vxlan (base border configurations) generated by template ext_base_border_vrflite_11_1

```
switch configure terminal
switch(config)#
ip prefix-list default-route seq 5 permit 0.0.0.0/0 le 1
ip prefix-list host-route seq 5 permit 0.0.0.0/0 eq 32
route-map extcon-rmap-filter deny 10
    match ip address prefix-list default-route
route-map extcon-rmap-filter deny 20
    match ip address prefix-list host-route
route-map extcon-rmap-filter permit 1000
route-map extcon-rmap-filter-allow-host deny 10
    match ip address prefix-list default-route
route-map extcon-rmap-filter-allow-host permit 1000
ipv6 prefix-list default-route-v6 seq 5 permit 0::/0
ipv6 prefix-list host-route-v6 seq 5 permit 0::/0 eq 128
route-map extcon-rmap-filter-v6 deny 10
    match ipv6 address prefix-list default-route-v6
route-map extcon-rmap-filter-v6 deny 20
    match ip address prefix-list host-route-v6
route-map extcon-rmap-filter-v6 permit 1000
route-map extcon-rmap-filter-v6-allow-host deny 10
    match ipv6 address prefix-list default-route-v6
route-map extcon-rmap-filter-v6-allow-host permit 1000
```

### Border-Vxlan VRF Lite Extension configuration

```
switch configure terminal
vrf context CORP
  ip route 0.0.0.0/0 2.2.2.2
exit
```

```
router bgp 100
  vrf CORP
    address-family ipv4 unicast
      network 0.0.0.0/0
      exit
    neighbor 2.2.2.2
      remote-as 200
      address-family ipv4 unicast
        send-community both
        route-map extcon-rmap-filter out
configure terminal
interface ethernet1/1.2
  encapsulation dot1q 2
  mtu 9216
  vrf member CORP
  ip address 2.2.2.22/24
  no shutdown
configure terminal
```

### WAN-Vxlan (External fabric Edge Router) VRF Lite Extension configuration

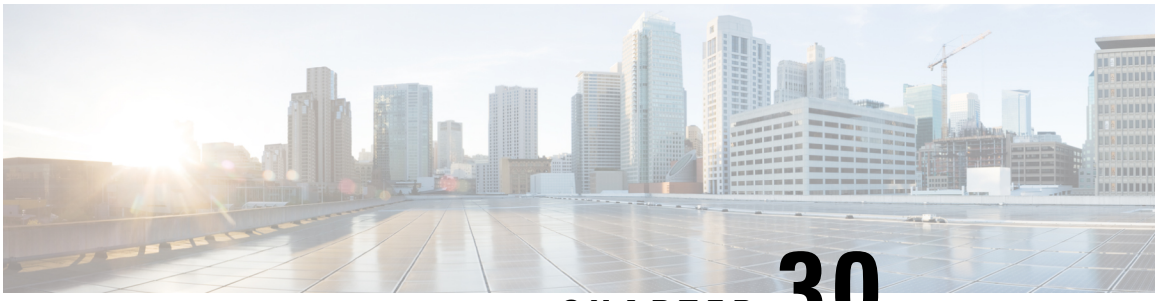
```
switch configure terminal
vrf context CORP
  address-family ipv4 unicast
exit
router bgp 200
  vrf CORP
    address-family ipv4 unicast
    neighbor 10.33.0.2
      remote-as 100
      address-family ipv4 unicast
        send-community both
        exit
      exit
    neighbor 10.33.0.6
      remote-as 100
      address-family ipv4 unicast
        send-community both
configure terminal
interface ethernet1/1.2
  mtu 9216
  vrf member CORP
  encapsulation dot1q 2
  ip address 10.33.0.1/30
  no shutdown
interface ethernet1/2.2
  vrf member CORP
  mtu 9216
  encapsulation dot1q 2
  ip address 10.33.0.5/30
  no shutdown
configure terminal
```



# PART IX

# Easy Provisioning of MSDC Deployments


- Managing eBGP Routed Fabrics, on page 647



CHAPTER 30

# Managing eBGP Routed Fabrics


- Managing BGP-Based Routed Fabrics, on page 647


# Managing BGP-Based Routed Fabrics

This chapter describes how to configure a typical spine-leaf based routed fabric with eBGP as the routing protocol of choice. This is the preferred deployment choice for Massively Scalable Data Center (MSDC) networks. Both Same-Tier-AS and Multi-AS options are supported. A routed fabric has no Layer-2 stretch or subnet stretch across leafs. In other words, networks are localized to a pair of leafs or a rack, with leafs hosting the default gateway for the directly attached server workloads. Subnet advertisement across racks are communicated over eBGP via the spine thereby providing any-to-any reachability within the routed fabric. Routed Fabric can be IPv4 or IPv6 based. IPv6 routed fabric uses IPv6 to build the intra-fabric connectivity and route advertisement. IPv6 routed fabric assigns link local address for intra-fabric links and support RFC 5549 to allow IPv4 route advertising using IPv6 next hop. Switch roles leaf, spine, border, super spine, and border super spine are supported.

## Creating an eBGP-based Fabric

1. Choose **LAN** > **Fabrics**.
2. From the **Actions** drop-down list, choose **Create Fabric**.

   The **Create Fabric** window appears.

   The fields are explained:

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

   **Fabric Template** - Click on this to choose the **BGP Fabric** fabric template. The fabric settings for creating a standalone fabric appear. Click **Select**.
3. The **General Parameters** tab is displayed by default.
4. The fields in the **General Parameters** tab are:

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

   **BGP AS Mode**: Choose **Multi-AS** or **Same-Tier-AS**.

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

In a **Same-Tier-AS** fabric, the spine switches have a unique BGP AS number and the leaf switches have a unique AS number, the borders share one AS. Leaf or border switches with the same role cannot have different AS.

Leafs and borders can have the same AS, or different AS.

The fabric is identified by the spine switch AS number.

**Manual Underlay IP Address Allocation**: Check the **Manual Underlay IP Address Allocation** check box to disable Dynamic Underlay IP Address Allocations.

5. Click **EVPN**. The Enable EVPN VXLAN Overlay option must be explicitly disabled. Note that this check box is enabled by default. This option should be enabled only for use-cases where customers want to build an eBGP-underlay/overlay based VXLAN EVPN fabric.

   **Routed Fabric**: In a Routed Fabric, once the IP reachability between the spine-leaf network has been established, you can easily create and deploy networks on the leafs using either HSRP or VRRP as the First-Hop Routing Protocol (FHRP) of choice.

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

   Note that **Routed_Network_Universal Template** is only applicable to a Routed Fabric.

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



**Note**

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

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

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

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

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

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

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

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

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

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

   You can enable the advertise PIP feature on a specific vPC as well.

**Enable the same vPC Domain Id for all vPC Pairs**: Check the **Enable the same vPC Domain Id for all vPC Pairs** check box. 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.



**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**.

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

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

**BGP Maximum Paths** - Specifies the BGP maximum paths.

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

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

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



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

**Enable BFD**: Check the **Enable BFD** 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.

NDFC supports BFD within a fabric. 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 through per switch freeform or per interface freeform policies.

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

```
feature bfd
```



**Note** NDFC with BFD enabled, 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 NDFC*.

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

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

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

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

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

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

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

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

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

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

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

**Enable CDP for Bootstrapped Switch**: Check the **Enable CDP for Bootstrapped Switch** check box to enable CDP for bootstrapped switch.

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

**Enable NX-API on HTTP**: Specifies enabling of NX-API on HTTP. Check **Enable NX-API on HTTP** and **Enable NX-API** check boxes 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 NDFC, 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 **Enable NX-API** and **Enable NX-API on HTTP** check boxes, applications use HTTP.

**Enable Strict Config Compliance**: Check the **Enable Strict Config Compliance** check box to enable this feature.

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



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

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

**Enable DCNM as Trap Host**: Check the **Enable DCNM as Trap Host** check box to enable NDFC as a trap host.

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

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

**Enable Default Queuing Policies**: Check **Enable Default Queuing Policies** 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. Pre-defined QoS configurations are included that can be used for various Cisco Nexus 9000 Series Switches. When you check this check box, the appropriate QoS configurations are pushed to the switches in the fabric. The system queuing is updated when configurations are deployed to the switches. You can perform the interface marking with defined queuing policies, if required, by adding the required configuration to the per interface freeform block.

Review the actual queuing policies by opening the policy file in the template editor. From Cisco NDFC Web UI, choose **Operations** > **Template**. 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 99.

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

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

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

9. Click **Manageability**. The fields in this tab are:

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

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

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

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

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

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

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

**AAA Freeform Config**: Specifies the AAA freeform configs.

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

10. Click **Bootstrap** tab. The fields in this tab are:

**Enable Bootstrap**: Check the **Enable Bootstrap** 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: Check the **Local DHCP Server** check box and enter details for the remaining mandatory fields.

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

If you do not check this check box, NDFC 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 NDFC IPv6 POAP is not supported with Cisco Nexus 7000 Series Switches. Cisco Nexus 9000 and 3000 Series Switches support IPv6 POAP only when switches are either L2 adjacent (eth1 or out-of-band subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other than /64 are not supported.

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

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

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

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

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

**Enable AAA Config**: Check the **Enable AAA Config** check box to include AAA configs from the **Manageability** tab during device bootup.

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

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

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

The format of the scope should be defined as:

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

For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

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

    **Hourly Fabric Backup**: Check the **Hourly Fabric Backup** 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, NDFC takes a backup.

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

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

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

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

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

**Note** Hourly and scheduled backup processes happen only during the next periodic configuration compliance activity, and there can be a delay of up to an hour. To trigger an immediate backup, do the following:

**a.** Choose **LAN** > **Topology**.

**b.** Click within the specific fabric box. The fabric topology screen comes up.

**c.** From the **Actions** pane at the left part of the screen, click **Re-Sync Fabric**.

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

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

**12.** Click **Flow Monitor**. The fields on this tab are:

**Enable Netflow** – Check the **Enable Netflow** check box to enable Netflow on VTEPs for this Fabric. By default, Netflow is disabled. On Enable, NetFlow configuration will be applied to all VTEPS that support netflow.

**Note** When Netflow is enabled on the fabric, you can choose not to have netflow on a particular switch by having a dummy no_netflow PTI.

If netflow is not enabled at the fabric level, an error message is generated when you enable netflow at the interface, network, or vrf level. For information about Netflow support for Cisco NDFC, refer Netflow Support, on page 149.

In the **Netflow Exporter** area, click **Actions** > **Add** to add one or more Netflow exporters. This exporter is the receiver of the netflow data. The fields on this tab are:

- **Exporter Name**: Specifies the name of the exporter.
- **IP**: Specifies the IP address of the exporter.
- **VRF**: Specifies the VRF over which the exporter is routed.
- **Source Interface**: Enter the source interface name.
- **UDP Port**: Specifies the UDP port over which the netflow data is exported.

Click **Save** to configure the exporter. Click **Cancel** to discard. You can also choose an existing exporter and select **Actions** > **Edit** or **Actions** > **Delete** to perform relevant actions.

In the **Netflow Record** area, click **Actions** > **Add** to add one or more Netflow records. The fields on this screen are:

- **Record Name**: Specifies the name of the record.
- **Record Template**: Specifies the template for the record. Enter one of the record templates names. In Release 12.0.2, the following two record templates are available for use. You can create custom netflow record templates. Custom record templates saved in the template library are available for use here.
  - **netflow_ipv4_record**: to use the IPv4 record template.

- **netflow_l2_record**: to use the Layer 2 record template.

- **Is Layer2 Record**: Check the **Is Layer2 Record** check box if the record is for Layer2 netflow.

Click **Save** to configure the report. Click **Cancel** to discard. You can also choose an existing record and select **Actions** > **Edit** or **Actions** > **Delete** to perform relevant actions.

In the **Netflow Monitor** area, click **Actions** > **Add** to add one or more Netflow monitors. The fields on this screen are:

- **Monitor Name**: Specifies the name of the monitor.
- **Record Name**: Specifies the name of the record for the monitor.
- **Exporter1 Name**: Specifies the name of the exporter for the netflow monitor.
- **Exporter2 Name**: (optional) Specifies the name of the secondary exporter for the netflow monitor.

The record name and exporters referred to in each netflow monitor must be defined in "**Netflow Record**" and "**Netflow Exporter**".

Click **Save** to configure the monitor. Click **Cancel** to discard. You can also choose an existing monitor and select **Actions** > **Edit** or **Actions** > **Delete** to perform relevant actions.

13. Click on the **Fabric** to view summary in the slide-in pane. Click on the Launch icon to view the Fabric Overview.

### Salient Points

- Brownfield migration is not supported for eBGP fabrics.
- You cannot change the leaf switch AS number after it is created and **Recalculate & Deploy** operation is executed. You need to delete the **leaf_bgp_asn** policy and execute **Recalculate & Deploy** operation to remove BGP configuration related to this AS first. Then, you can add the leaf_bgp_asn policy with the new AS number.
- If you want to switch between Multi-AS and Same-Tier-AS modes, remove all manually added BGP policies (including leaf_bgp_asn on the leaf switch and the ebgp overlay policies), and execute the **Recalculate & Deploy** operation before the mode change.
- The supported roles are leaf, spine, super spine, border leaf, and border super spine.
- On the border and super spine border device, VRF-Lite is supported with manual mode

# Adding Switches to a Fabric

Switches in each fabric are unique, and hence, each switch can only be added to one fabric. see Adding Switches to a Fabric, on page 279.

# Deploying Fabric Underlay eBGP Policies

In NDFC, a fabric with the **BGP Fabric** template is created. One spine switch and three leaf switches are imported to it.

The two different types of fabrics are:

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

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

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

To add a policy to the required switch, see Adding a Policy, on page 195.

# Deploying Networks in eBGP-based Fabrics

## Overview of Networks in a Routed Fabric

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



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

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

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

### Guidelines

- HSRP authentication or VRRP authentication is not supported. If you want to use authentication, you can enter the applicable commands in the network freeform config.
- vPC peer gateway can be used to minimize peer link usage in the case that some third-party devices ignore the HSRP virtual-MAC and use the ARP packet source MAC for ARP learning. In Routed fabric mode, NDFC generates vPC peer gateway command for VPC devices.
- For an eBGP fabric, changing between routed fabric type and EVPN fabric type, or HSRP and VRRP, is not allowed with the presence of networks and VRFs. You need to undeploy and delete these networks

and VRFs before changing the fabric type or FHRP. For more information, see *Undeploying Networks for the Standalone Fabric* and *Undeploying VRFs for the Standalone Fabric*.

- If the fabric was running in Routed Fabric mode previously, the default fabric values such as FHRP protocol and network VLAN range are internally set for a Routed Fabric. You need to edit the fabric settings if you want to configure different values. Before deploying a network configuration, you need to update the FHRP protocol fabric setting and click **Recalculate & Deploy**.

## Creating and Deploying a Network in a Routed Fabric

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

### Before you begin

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

### Procedure

**Step 1** Choose any one of the following navigation paths:

- Choose **LAN > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Networks**.
- Choose **LAN > Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Networks**.

**Step 2** From the **Actions** drop-down list, choose **Create**.

The **Create Networks** window appears. The fields in this window are:

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

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

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

**Network Template**: Select the **Routed_Network_Universal** template.

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

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

In the **General Parameters**tab, specify the required details.

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

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

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

**Interface IPv6 addr on active**: Specifies the IPv6 interface address on an active device in a vPC pair. This field is applicable only when you are creating and deploying a network for a vPC pair of devices.

**Interface IPv6 addr on standby**: Specifies the IPv6 interface address on a standby/backup device in a vPC pair. This field is applicable only when you are creating and deploying a network for a vPC pair of devices.

**IPv6 Link Local address**: Specifies the IPv6 link local address. This field is applicable only when you are creating and deploying a network for a vPC pair of devices and VRRP is chosen as the FHRP protocol.

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

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

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

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

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

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

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

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

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

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

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

**Enable Preempt**: Specifies whether the standby/backup device can preempt an active device.

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

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

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

**Step 3** Click **Create Network**. For more information, see Networks, on page 211.

**Step 4** In the **Network Attachment** window, for a vPC pair, assign the active state for a device.

Check the **isActive** check box for an active device and uncheck the **isActive** check box for a standby device.

Click **Save**.

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

**Step 5** (Optional) Click the **Preview** icon to preview the configuration that deployed on devices.

The **Preview Configuration** window is displayed.

**Step 6** Click **Deploy**.

You can also deploy the network by navigating to the **Fabric Overview** window and clicking the **Deploy** button.

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

You can use an inter-fabric link to connect a route fabric to an edge router. This link configures an IP address on the physical interface and establish eBGP peering with the edge router on default vrf. The BGP configuration includes advertising default route to leaf switches.

**Note** The **Fabric Monitor Mode** check box in the external fabric settings can be unchecked. Unchecking the **Fabric Monitor Mode** check box enables NDFC to deploy configurations to the external fabric. For more information, see Creating an External Fabric, on page 112.

### Procedure

**Step 1** Choose **LAN > Fabrics**. Double-click on a routed fabric.

The **Fabric Overview** window appears.

**Step 2** On the **Links** tab, click **Actions > Create**.

The **Link Management - Create Link** window appears.

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

**Link Sub-Type**: This field populates the IFC type. Choose **ROUTED_FABRIC** from the drop-down list.

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

**Source Fabric**: This field is pre-populated with the source fabric name.

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

**Source Device** and **Source Interface**: Choose the source device and Ethernet or port channel interface that connects to the destination device. Only device with the border role can be chosen.

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

Based on the selection of the source device and source interface, the destination information is auto populated 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.

The **General Parameters** tab contains the following fields.

**Source BGP ASN**: In this field, the AS number of the leaf is auto populated if you have created and applied the **leaf_bgp_asn** policy.

**Source IPv4 Address/Mask**: Fill up this field with the IP address of the source interface that connects to the destination device.

**Destination IPv4**: Fill up this field with the IPv4 address of destination interface

**Destination BGP ASN**: In this field, the AS number of the destination device is auto populated.

**Source IPv6 Address/Mask**: Fill up this field with the IP address of the source interface that connects to the destination device.

**Destination IPv6**: Fill up this field with the IPv6 address of destination interface.

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

**Link MTU**: Fill up this field with the interface MTU.

**Disable Default Route Config**: Check the **Disable Default Route Config** check box.

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

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

**Source Interface Freeform CLIs** and **Destination Interface Freeform CLIs**: Enter the freeform configurations specific to the source and destination interfaces. You should add the configurations as displayed in the running configuration of the switch, without indentation. For more information, refer Enabling Freeform Configurations on Fabric Switches , on page 95.

**Step 3** Click **Save**.

**Step 4** Double-click the device which is connecting to the edge router in the external fabric, and click **Actions > Recalculate & Deploy**.

**Step 5** After the configuration deployment is completed in the **Deploy Configuration** window, click **Close**.

**Step 6** Navigate to the external fabric in the **LAN Fabric** window, and double-click on the fabric.

**Step 7** Click the **Links** tab to see all the links for the external fabric.

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

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

**Step 8** Navigate to the **LAN Fabric** window.

**Step 9** Double-click the external fabric connecting to the routed fabric and click **Actions > Recalculate & Deploy**.

**Step 10** After the configuration deployment is completed in the **Deploy Configuration** window, click **Close**.