



Cisco Nexus Dashboard Fabric Controller for LAN Configuration Guide, Release 12.0.1a

First Published: 2021-09-30

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 NON-INFRINGEMENT 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 Cisco Systems, Inc. All rights reserved.



CHAPTER 1

Overview

- [Know your Web UI, on page 1](#)
- [User Feedback, on page 2](#)

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:

- **Cisco Nexus Dashboard Fabric Controller Personality** - Specifies the deployment personality.
- **Feedback** - Allows you to provide feedback on Cisco Nexus Dashboard Fabric Controller. For instructions, see [User Feedback, on page 2](#).
- **Help** - Click on the **Help** icon 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.
 - **API Documentation** - Click to launch the API documentation.
- **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:

- **Logout** - Allows you to terminate the Web UI and return to the login screen.
- **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.
- **User Preferences** - Allows you to view the Welcome screen on every login.

- **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:
- On the Cisco Nexus Dashboard Web UI, choose **Infrastructure > Cluster Configuration**.
The **Cluster Configuration General** tab appears.
 - In the **Proxy Configuration** area, click **Edit** icon.
 - In the **Servers** area, click **Add Server**.
 - Choose the protocol type – **HTTP** or **HTTPS**.
 - 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.
-



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 5](#)
- [Viewing vCenter VMs, on page 6](#)
- [Viewing Kubernetes Pods, on page 6](#)
- [Monitoring Endpoint Locator, on page 7](#)

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 summary of problems, and number in the donut depicting total number of fabrics. Displays fabric health status with Critical , and Healthy .
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 Critical , and Healthy with the corresponding count.
Switch Roles	Displays the switches roles summary and the corresponding count. Displays the number of access, spine and leaf devices.

Dashlet	Description
Switch Hardware Version	Displays the switches models and the corresponding count.
Switch Software Version	Displays the switches software version and the corresponding count.
Reports	Displays switch reports.

Viewing vCenter VMs



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.

For more information regarding this feature, refer to Cisco Nexus Dashboard Fabric Controller for LAN Configuration Guide.

UI Path: **Dashboard > vCenter VMs**

The vCenter VMs tab displays the following details of VMs:

- VM Name, its IP address and MAC address
- Host name that is connected to a VM
- Switch name that is connected to a VM, switch's IP address, MAC address, and interface
- Port channel ID and VPC ID
- VLAN segment type
- Power state of the VM
- Physical NIC of VM

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

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.

For more information regarding this feature, refer to Cisco Nexus Dashboard Fabric Controller for LAN Configuration Guide.

UI Path: **Dashboard > 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 Container Orchestrator 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 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 pod.
Physical Switch	Specifies the physical switch connected to pod.
Switch Interface	Specifies the switch interface connected to pod.
Cluster Name	Specifies the name of the cluster.
Port Channel	Specifies the port channel.
VLAN	Specifies the VLAN.
Fabric	Specifies the fabric name.

Monitoring Endpoint Locator

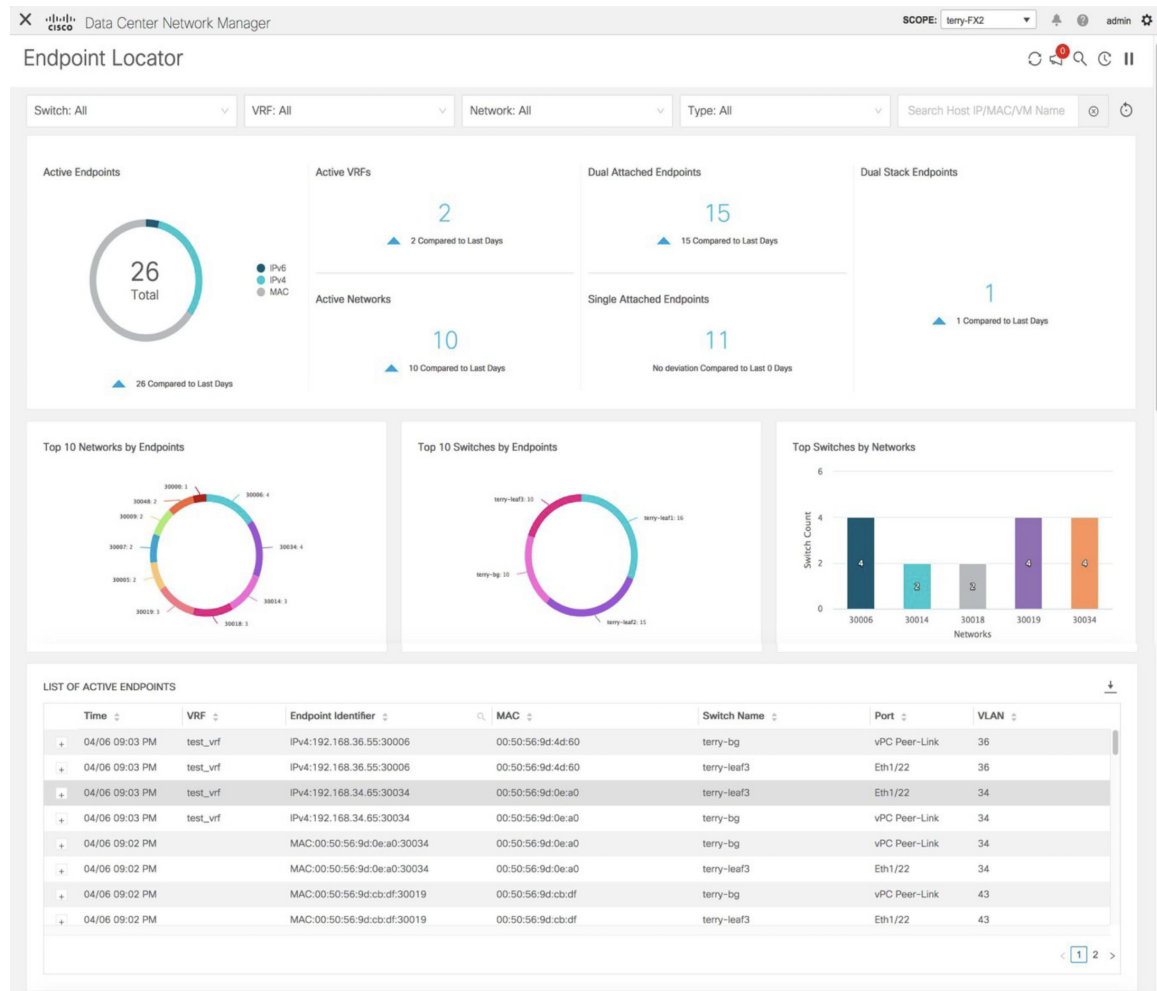
Information about the Endpoint Locator is displayed on a single landing page or dashboard. The dashboard displays an almost real-time view of data (refreshed every 30 seconds) pertaining to all the active endpoints on a single pane. The data that is displayed on this dashboard depends on the scope selected 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.

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.



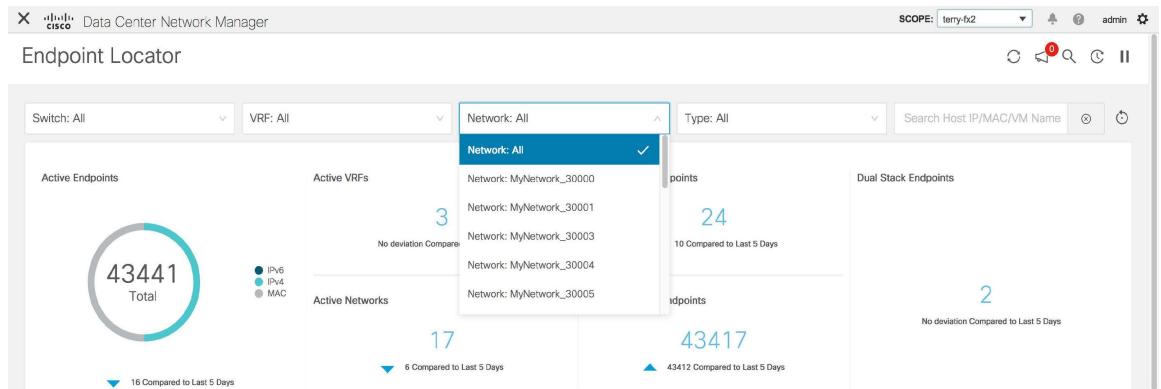
Note

Due to an increase in scale from Cisco Nexus Dashboard Fabric Controller Release 11.3(1), the system may take some time to collect endpoint data and display it on the dashboard. Also, 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 also filter and view the endpoint locator details for a specific **Switch**, **VRF**, **Network**, and **Type**, by using the respective drop-down lists. Starting from Cisco Nexus Dashboard Fabric Controller Release 11.3(1), 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 also 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.



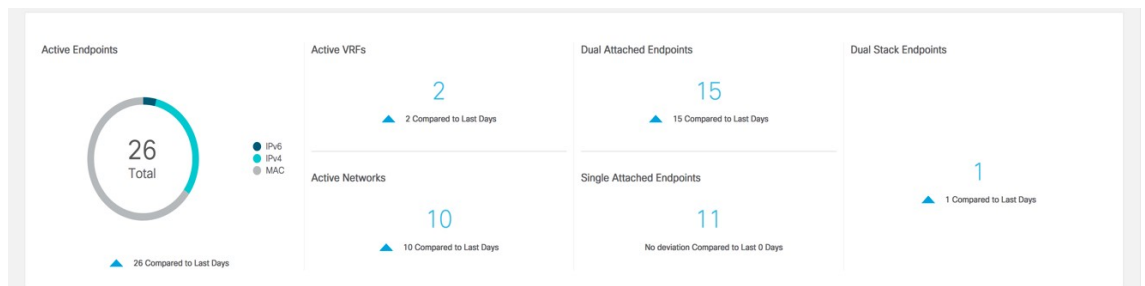
Note You can initiate a search by using the available options from the dropdown lists or by using the **Search Host IP/MAC/VM Name** field. You cannot initiate a search by using a combination of dropdown lists and the search field.



You can reset the filters to the default options by clicking the **Reset Filters** icon.



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.

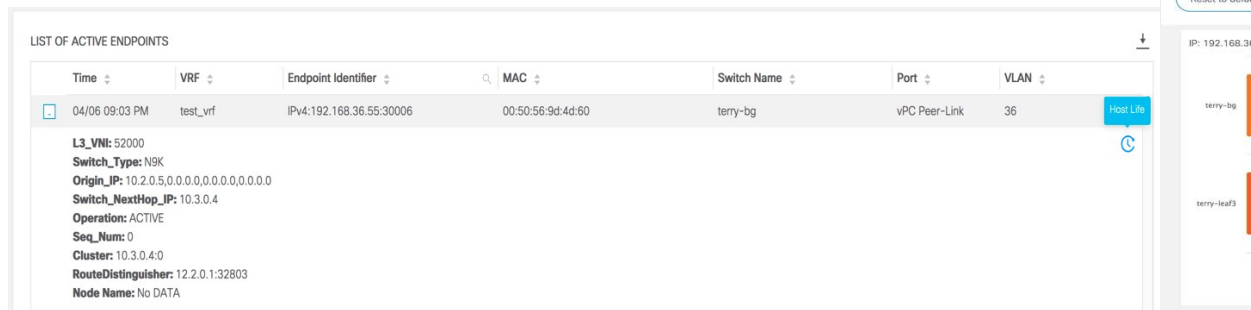
Time	VRF	Endpoint Identifier	MAC	Switch Name	Port	VLAN
04/06 09:03 PM	test_vrf	IPv4:192.168.36.55:30006	00:50:56:9d:4d:60	terry-bg	vPC Peer-Link	36
04/06 09:03 PM	test_vrf	IPv4:192.168.36.55:30006	00:50:56:9d:4d:60	terry-leaf3	Eth1/22	36
04/06 09:03 PM	test_vrf	IPv4:192.168.34.65:30034	00:50:56:9d:0e:a0	terry-leaf3	Eth1/22	34
04/06 09:03 PM	test_vrf	IPv4:192.168.34.65:30034	00:50:56:9d:0e:a0	terry-bg	vPC Peer-Link	34
04/06 09:02 PM		MAC:00:50:56:9d:0e:a0:30034	00:50:56:9d:0e:a0	terry-bg	vPC Peer-Link	34
04/06 09:02 PM		MAC:00:50:56:9d:0e:a0:30034	00:50:56:9d:0e:a0	terry-leaf3	Eth1/22	34
04/06 09:02 PM		MAC:00:50:56:9d:cb:df:30019	00:50:56:9d:cb:df	terry-bg	vPC Peer-Link	43
04/06 09:02 PM		MAC:00:50:56:9d:cb:df:30019	00:50:56:9d:cb:df	terry-leaf3	Eth1/22	43

Click + to display more information for a specific endpoint. 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.

Time	VRF	Endpoint Identifier	MAC	Switch Name	Port	VLAN
06/11 09:39 AM	myvrf_50001	IPv6:2188:1::99:30001	00:50:56:be:71:e9	leg-fab2-bgw2	Po606	2344

L3_VNI: 50001
Switch_Type: NGK
Origin_IP: 40.4.0.1,0.0.0.0,0.0.0.0,0.0.0.0
Switch_NextHop_IP: 40.3.0.2
Operation: ACTIVE
Seq_Num: 0
Cluster: 40.3.0.2.0
RouteDistinguisher: 40.2.0.1:35111
Node Name: ppp-leg-fab2-188

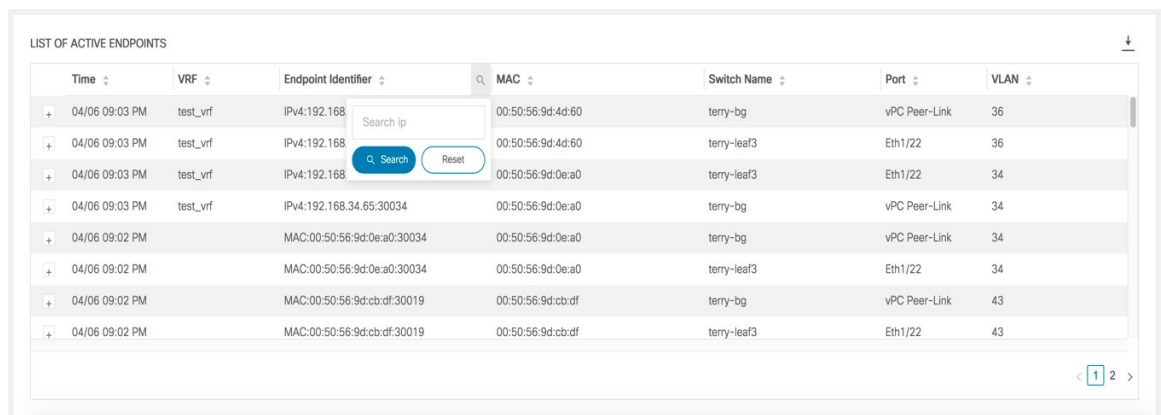
Click the **Host Life** icon to display the **Endpoint Life** window for that endpoint.



Time	VRF	Endpoint Identifier	MAC	Switch Name	Port	VLAN
04/06 09:03 PM	test_vrf	IPv4:192.168.36.55:30006	00:50:56:9d:4d:60	terry-bg	vPC Peer-Link	36

L3_VNI: 52000
Switch_Type: N9K
Origin_IP: 10.2.0.5,0.0.0.0,0.0.0.0,0.0.0.0
Switch_NextHop_IP: 10.3.0.4
Operation: ACTIVE
Seq_Num: 0
Cluster: 10.3.0.4:0
RouteDistinguisher: 12.2.0.1:32803
Node Name: No DATA


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




Time	VRF	Endpoint Identifier	MAC	Switch Name	Port	VLAN
04/06 09:03 PM	test_vrf	IPv4:192.168.36.55:30006	00:50:56:9d:4d:60	terry-bg	vPC Peer-Link	36
04/06 09:03 PM	test_vrf	IPv4:192.168.36.55:30006	00:50:56:9d:4d:60	terry-leaf3	Eth1/22	36
04/06 09:03 PM	test_vrf	IPv4:192.168.36.55:30006	00:50:56:9d:0e:a0	terry-leaf3	Eth1/22	34
04/06 09:03 PM	test_vrf	IPv4:192.168.34.65:30034	00:50:56:9d:0e:a0	terry-bg	vPC Peer-Link	34
04/06 09:02 PM		MAC:00:50:56:9d:0e:a0:30034	00:50:56:9d:0e:a0	terry-bg	vPC Peer-Link	34
04/06 09:02 PM		MAC:00:50:56:9d:0e:a0:30034	00:50:56:9d:0e:a0	terry-leaf3	Eth1/22	34
04/06 09:02 PM		MAC:00:50:56:9d:cb:df:30019	00:50:56:9d:cb:df	terry-bg	vPC Peer-Link	43
04/06 09:02 PM		MAC:00:50:56:9d:cb:df:30019	00:50:56:9d:cb:df	terry-leaf3	Eth1/22	43

In certain scenarios, the datapoint database may go out-of-sync and information, such as the number of endpoints, may not be 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.

In such cases, clicking the **Resync**  icon leads to the dashboard syncing to the data currently in the RR. However, historical data is preserved. We recommend not clicking **Resync** multiple times as this is a compute-intensive activity.

Endpoint Locator



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

Notifications


Time

06/11 05:30 AM

06/11 05:28 AM


06/10 07:03 AM

Endpoint Locator




Information such as the time at which the notification was generated, the description of the notification, severity level, and the name of the node 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. Click the download icon to download the list of notifications as a CSV file.

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.

Endpoint Locator



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 the **Resync** icon at the top right of the EPL dashboard.

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 180 days amounting to a maximum of 100 GB storage space.



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 hover over the 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 the download icon at the top right of the **Endpoints** window to download the data as a CSV file.


Endpoints ↓ ×

Endpoint	Switch Name	VRF
IPv4:192.168.36.20:30006	terry-leaf3	test_vrf
IPv4:192.168.200.2:32000	terry-leaf3	test_vrf
IPv4:192.168.36.29:30006	terry-leaf2	test_vrf
IPv4:192.60.0.100:30004	terry-leaf1	myvrf_50000
IPv4:192.168.80.90:30080	terry-leaf1	test_vrf
IPv4:192.168.180.100:30008	terry-leaf3	myvrf_50009
IPv4:192.168.48.2:30048	terry-leaf2	test_vrf
IPv4:192.168.39.2:30043	terry-leaf2	test_vrf
IPv4:192.60.7.208:30004	terry-leaf3	myvrf_50000
IPv4:192.60.10.168:30004	terry-leaf3	myvrf_50000

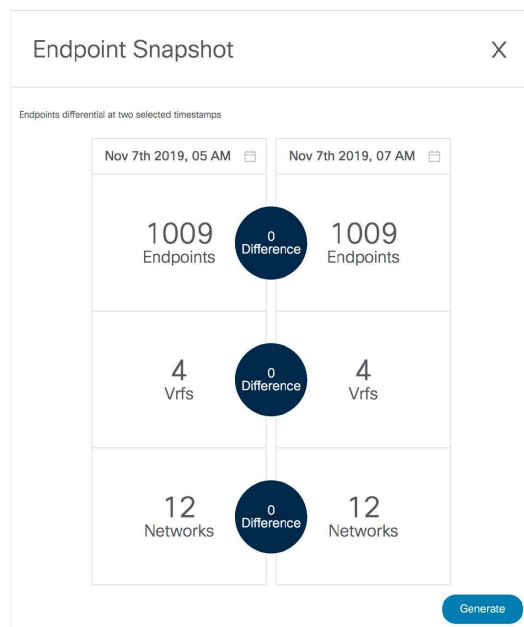
< 1 2 3 4 5 ... 303 >

Endpoint Snapshots


You can compare endpoint data at two specific points in time. To display the **Endpoint Snapshot** window, click the **Endpoint Snapshot** icon at the top right of the **Active Endpoints** graph in the **Endpoint History** window.




By default, endpoint snapshot comparison data for the previous hour is displayed.



To compare endpoint snapshots at specific points in time, select two points in time, say T1 and T2, and click **Generate**.




A comparison of the endpoints, VRFs, and networks at the selected points in time are displayed. Click each tile to download more information about the endpoints, VRFs, or networks. Click the **Difference** icon to download details about the differences in data for the specified time interval. Snapshots are stored for a maximum of three months and then discarded.

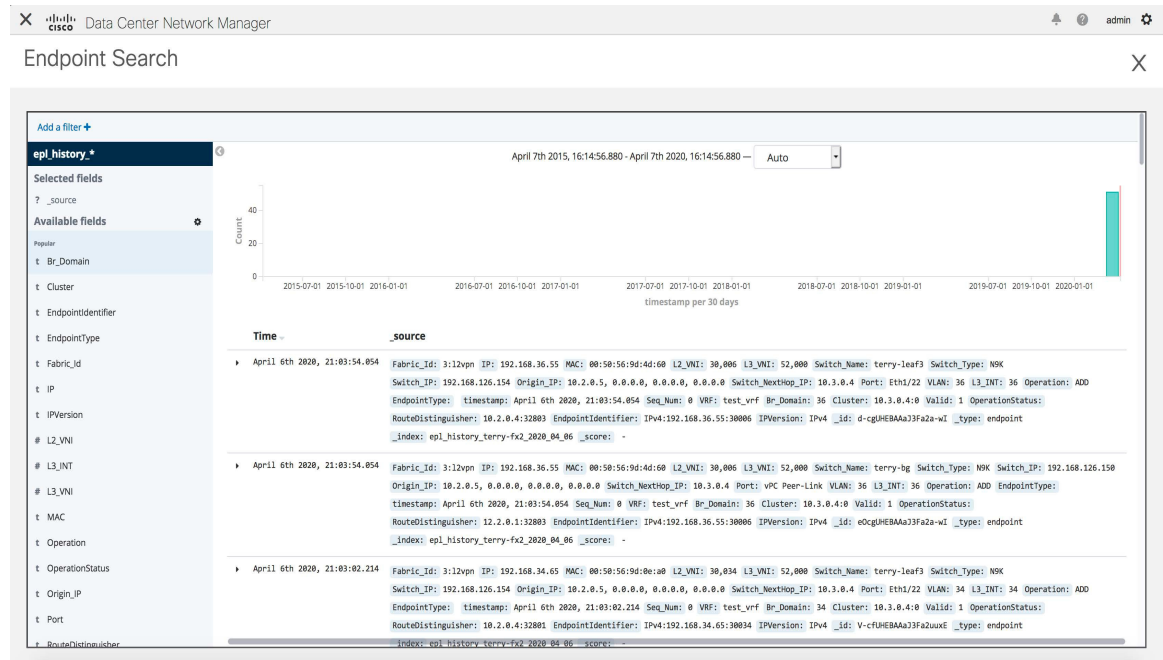


Endpoint Search

Click the **Endpoint Search** icon at the top right of the Endpoint Locator landing page to view a real-time plot displaying endpoint events for the period specified in a date range.




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 fields.



Endpoint Life

Click the **Endpoint Life** icon at the top right of the Endpoint Locator landing page to display a time line of a particular endpoint in its entire existence within the fabric.



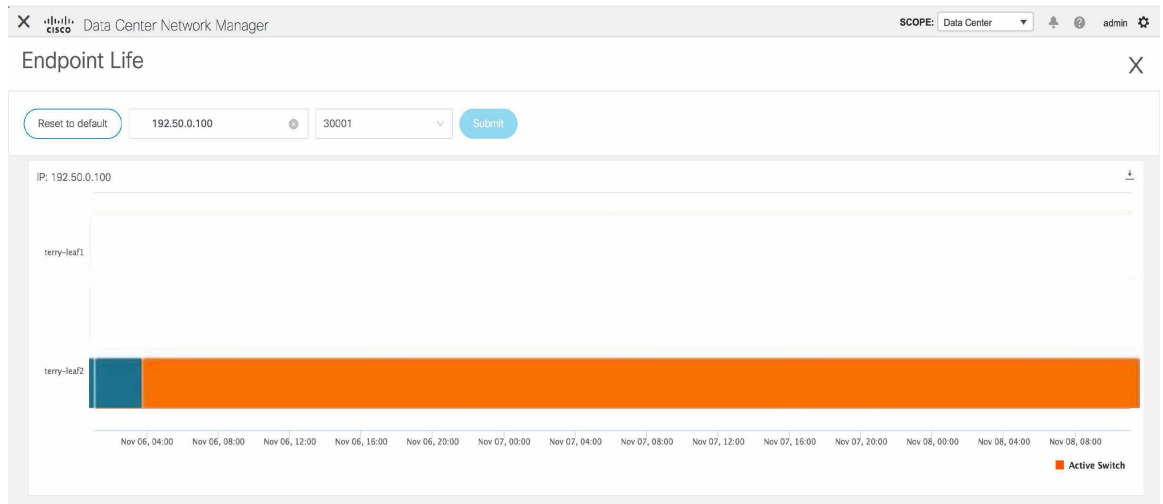
Specify the IP or MAC address of an endpoint and the VXLAN Network Identifier (VNI) to display the list of switches that an endpoint was present under, including the associated start and end dates. Click **Submit**.

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 screenshot displays the 'Endpoint Life' interface in Cisco Data Center Network Manager. It shows a search bar with a 'SCOPE' dropdown set to 'terry-fx2'. Below the search bar, there is a 'Reset to default' button, an input field for 'Enter IP or MAC', a 'Select VNI' dropdown, and a 'Submit' button. Below the search bar, there is a message: 'Please enter IP & VNI to see the graph'.

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 window.





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 119](#).
- 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-pollled CPU utilization, and last-pollled 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 221](#).
- 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 117](#).

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 20](#)
- [Viewing Topology, on page 21](#)

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.

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 22](#).
- 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 22](#).
- 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 23](#).



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.

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.
- After moving a cable from one port to another port, the old link is retained in the **Topology** window, and it's shown in the red color indicating that the link is down. The port movements aren't updated in the **Topology** window. Rediscover the switch for the updated ports to be displayed in Nexus Dashboard Fabric Controller.

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

Virtual Management

- [Virtual Infrastructure Manager, on page 27](#)



CHAPTER 4

Virtual Infrastructure Manager

- [Virtual Infrastructure Manager, on page 27](#)
- [Enabling Virtual Infrastructure Manager, on page 28](#)
- [Enabling Kubernetes Cluster, on page 29](#)

Virtual Infrastructure Manager

UI Path: **Virtual Management** > **Virtual Infrastructure Manager**

In virtualized environments, any kind of troubleshooting starts with identifying the network attachment point for the virtual machines. This means that a quick determination of the server, virtual switch, port group, VLAN, associated network switch, and physical port is critical. This requires multiple touch points and interactions between the server and the network administrator as well as reference to multiple tools (compute orchestrator, compute manager, network manager, network controller, and so on).

This allows you to visualize the vCenter-managed hosts and their leaf switch connections on the Dashboard **vCenter VMs** tab. The visualization options include viewing only the attached physical hosts, only the VMs, or both.



Note

- The vCenter Compute Visualization feature is supported on both the LAN Classic and Easy Fabrics installations for the vCenter-managed computes.
- It is not recommended to use special characters in a VM name as vCenter does not escape special characters used in display names. For more information, see <https://vss-wiki.eis.utoronto.ca/display/VSSPublic/Virtual+Machine+Naming>.
- Cisco Nexus Dashboard Fabric Controller doesn't support non-Cisco blade servers.

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
Managed	Specifies the status either Managed or Unmanaged.

Field	Description
Status	Specifies the status of vCenter.
User	Specifies the user created the vCenter.
Last Updated Time	Specifies the last updated time for vCenter.

Enabling Virtual Infrastructure Manager



Note Ensure that you have enabled Network visualization of Virtual Machines feature for Cisco Nexus Dashboard Fabric Controller .

Choose **Settings > Feature Management** choose check box **VMM Visualizer**.

You can view the added Virtual Infrastructure Manager details on dashboard. Navigate **Dashboard > vCenter VMs**

You can perform various actions with Virtual Infrastructure Manager feature from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps.

Procedure

Step 1 Choose **Virtual Management > Virtual Infrastructure Manager**.

The **Virtual Infrastructure Manager** window appears.

Step 2 Choose **Actions > Add Instance**

The **Add Instance** window appears.

- a) Choose **vCenter** from Select Type drop-down list.
- b) Enter required IP address and password in the respective fields.
- c) Click **Add**.

You can view added instance in the window.

Step 3 Choose required vCenter, choose **Actions > Edit Instance** to edit instance.

The **Edit Instance** window appears.

You can update password for the selected edit instance and also change the status to Managed or Unmanaged and click **Edit**.

Step 4 Choose required vCenter, choose **Actions > Delete Instance(s)** to delete the vCenter.

Step 5 Choose required vCenter, choose **Actions > Rediscover Instance(s)** to rediscover the vCenter.

The **Confirmation window** appears, click **Confirm**.

Step 6 Choose **Actions** > **Refresh** to refresh the Virtual Infrastructure Manager table.

Enabling 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**.

You can view the added Kubernetes Visualizer details on dashboard. Navigate **Dashboard** > **Kubernetes Pods**

You can perform various actions with Kubernetes Cluster feature from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps.

Procedure

Step 1 Choose **Virtual Management** > **Virtual Infrastructure Manager**.

The **Virtual Infrastructure Manager** window appears.

Step 2 Choose **Actions** > **Add Instance**

The **Add Instance** window appears.

- a) Choose **Kubernetes Cluster** from Select Type drop-down list.
- b) Enter Cluster IP address, Username, Cluster Certificate and Client Certificate in the respective fields.
- c) Click **Add**.

You can view added instance in the window.

Step 3 Choose required vCenter, choose **Actions** > **Edit Instance** to edit instance.

The **Edit Instance** window appears.

You can update password for the selected edit instance and also change the status to Managed or Unmanaged and click **Edit**.

Step 4 Choose required vCenter, choose **Actions** > **Delete Instance(s)** to delete the vCenter.

Step 5 Choose required vCenter, choose **Actions** > **Rediscover Instance(s)** to rediscover the vCenter.

The **Confirmation window** appears, click **Confirm**.

Step 6 Choose **Actions** > **Refresh** to refresh the Virtual Infrastructure Manager table.



PART II

LAN

- [Fabrics, on page 33](#)
- [Switches, on page 221](#)
- [Policies, on page 251](#)
- [Interfaces, on page 255](#)
- [Services, on page 271](#)



CHAPTER 5

Fabrics

- [LAN Fabrics](#), on page 33
- [Enhanced Role-based Access Control](#), on page 112
- [Nexus Dashboard Security Domains](#), on page 114
- [Backup Fabric](#), on page 115
- [Restoring Fabric](#), on page 117
- [VXLAN OAM](#), on page 117
- [Fabric Overview](#), on page 119
- [Endpoint Locator](#), on page 204

LAN Fabrics

The following terms are referred to in the document:

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

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.
Fabric Type	Displays the type of the fabric—Switch Fabric, LAN Monitor, or External
ASN	Displays the ASN for the fabric.

Field	Description
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 35 .
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.

This section contains the following topics:

Fabric Summary

Click on the 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 - This section provides basic about the Fabric.

Inventory - This section provides information about Switch Configuration and Switch Health.

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

Prerequisites to Create 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 Nexus Dashboard. For more information, refer to *Cluster Configuration* section in [Cisco Nexus Dashboard User Guide](#).

Override ESXi Networking for Promiscuous Mode

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 > Fabrics**.
- Step 2** From the **Actions** drop-down list, select **Create Fabric**.
- Step 3** Enter the fabric name and click **Choose Template**.
- Step 4** Based on your fabric requirements, select one of the fabric templates and click **Select**.
- Step 5** Specify the values for the fabric settings and click **Save**.
-

VXLAN BGP EVPN Fabrics Provisioning

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

Border devices in a fabric typically provide external connectivity via peering with appropriate edge/core/WAN routers. These edge/core routers may either be managed or monitored by 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.

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

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.

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

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

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

from the 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 its current running configuration, and a config compliance check is performed for the switch to get the **In-Sync** or **Out-of-Sync** status. If the switch is out of sync,

the user is provided with a preview of all the configurations running in that particular switch that vary from the intent defined by the user for that respective switch.

- Persistent configuration diff is seen for the command line: **system nve infra-vlan int force**. The persistent diff occurs if you have deployed this command via the freeform configuration to the switch. Although the switch requires the **force** keyword during deployment, the running configuration that is obtained from the switch in 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 New VXLAN BGP EVPN Fabric

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

This procedure contains descriptions for the IPv4 underlay.

1. From Actions drop-down list, choose **Create Fabric**.

The **Create Fabric** window appears.

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

The fields are explained:

Fabric Name - Enter the name of the fabric.

Choose Template - Click on this to choose the **Easy_Fabric** fabric template. The fabric settings for creating a standalone fabric appear. Click **Select**,

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



Note If you'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), then browse through the Multi-Site Domain for VXLAN BGP EVPN Fabrics topic before member fabric creation.

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

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

Enable IPv6 Underlay: Enable the IPv6 underlay feature. .

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.

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

To deploy spine devices as RRs, 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.

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

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

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

An instance of the **rr_state** policy is applied on the spine switch if it'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.

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

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

- d. Click **Deploy Config** in the fabric topology window.

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

Anycast Gateway MAC : Specifies the anycast gateway MAC address.

Enable Performance Monitoring: Check the check box to enable performance monitoring.

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

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

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

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

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

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

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

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**.

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

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

When you choose ASM, the BiDir related fields aren't enabled. When you choose BiDir, the BiDir related fields are enabled.



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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

You can enable the advertise PIP feature on a specific vPC as well. .

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. .



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**.

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

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

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

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

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



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

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

OSPF Authentication Key ID - The Key ID is populated.

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



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

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

Enable IS-IS 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.



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

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



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

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

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



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

Enable PIM Hello Authentication – 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.

To retrieve PIM Hello Authentication Key, perform the following steps:

- a. SSH into the switch.
- b. On an unused switch interface, enable the following:

```
switch(config)# interface e1/32
switch(config-if)# ip pim hello-authentication ah-md5 pimHelloPassword
```

In this example, **pimHelloPassword** is the cleartext password that has been used.

- c. Enter the **show run interface** command to retrieve the PIM hello authentication key.

```
switch(config-if)# show run interface e1/32 | grep pim
ip pim sparse-mode
ip pim hello-authentication ah-md5 3 d34e6c5abc7fecf1caa3b588b09078e0
```

In this example, **d34e6c5abc7fecf1caa3b588b09078e0** is the PIM hello authentication key that should be specified in the fabric settings.

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

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

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

```
feature bfd
```

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

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

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

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

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

Here are the examples of the BFD global policies:

```
router ospf <ospf tag>
  bfd

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

ip pim bfd

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

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



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

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

BFD Authentication Key – Specifies the BFD authentication key.

For information about how to retrieve the BFD authentication parameters. .

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

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

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

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

The following fields can be used to specify different configurations:

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

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

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

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.

Overlay Mode – VRF/Network configuration using config-profile or CLI, default is config-profile. For more information, see [Overlay Mode, on page 57](#).

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

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

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

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

Power Supply Mode – Choose the appropriate power supply mode.

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

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

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

Auto_Net_VNI\$\$VNI\$\$_VLAN\$\$VLAN_ID\$\$. When you create networks, the name is generated according to the syntax you specify. The following table describes the variables in the syntax.

Variables	Description
\$\$VNI\$\$	Specifies the network VNI ID found in the switch configuration. This is a mandatory keyword required to create unique network names.
\$\$VLAN_ID\$\$	Specifies the VLAN ID associated with the network. VLAN ID is specific to switches, hence 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. We recommend not to use this unless the VLAN ID is consistent across the fabric for the VNI.

Variables	Description
<string>	This variable is optional and you can enter any number of alphanumeric characters that meet the network name guidelines.

Example overlay network name: Site_VNI12345_VLAN1234



Note Ignore this field for greenfield deployments. The **Brownfield Overlay Network Name Format** applies for the following brownfield imports:

- CLI-based overlays
- Configuration profile-based overlay

Enable CDP for Bootstrapped Switch – Enables CDP on management (mgmt0) interface for bootstrapped switch. By default, for bootstrapped switches, CDP is disabled on the mgmt0 interface.

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

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



Note The VXLAN OAM feature in Cisco Nexus Dashboard Fabric Controller is only supported on a single fabric or site.

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



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

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

Enable NX-API on HTTP on Port – Specifies enabling of NX-API on HTTP. Enable this check box and the **Enable NX-API** check box to use HTTP. This check box is checked by default. If you uncheck this check box, the applications that use NX-API and supported by Cisco 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.

Enable Policy-Based Routing (PBR) – Select this check box to enable routing of packets based on the specified policy. Starting with Cisco NX-OS Release 7.0(3)I7(1) and later releases, this feature works on Cisco Nexus 9000 Series switches with Nexus 9000 Cloud Scale (Tahoe) ASICs. This feature

is used along with the Layer 4-Layer 7 service workflow. For information on Layer 4-Layer 7 service, refer the *Layer 4-Layer 7 Service* chapter.

Enable Strict Config Compliance – Enable the Strict Config Compliance feature by selecting this check box. By default, this feature is disabled.

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.

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

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

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

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

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

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

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

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

Enable Default Queuing Policies – Check this check box to apply QoS policies on all the switches in this fabric. To remove the QoS policies that you applied on all the switches, uncheck this check box, update all the configurations to remove the references to the policies, and save and deploy. 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.

The DSCP mapping for QoS 5 has changed from 40 to 46 in the policy template.

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.

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 [Enabling MACsec](#).

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

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

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

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

7. Click the **Resources** tab.

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

- By default, Nexus Dashboard Fabric Controller allocates the underlay IP address resources (for loopbacks, fabric interfaces, etc) dynamically from the defined pools. If you select the check box, the allocation scheme switches to static, and some of the dynamic IP address range fields are disabled.
- For static allocation, the underlay IP address resources must be populated into the Resource Manager (RM) using REST APIs.
- The Underlay RP Loopback IP Range field stays enabled if BIDIR-PIM function is chosen for multicast replication.
- Changing from static to dynamic allocation keeps the current IP resource usage intact. Only future IP address allocation requests are taken from dynamic pools.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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



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

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

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

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

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 **AAA Configurations** will be created.

9. Click the **Bootstrap** tab.

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

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

- External DHCP Server: Enter information about the external DHCP server in the **Switch Mgmt Default Gateway** and **Switch Mgmt IP Subnet Prefix** fields.

- Local DHCP Server: Enable the **Local DHCP Server** check box and enter details for the remaining mandatory fields.

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

If you do not select this check box, 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.



Note Cisco Nexus 9000 and 3000 Series Switches support IPv6 POAP only when switches are either Layer-2 adjacent (eth1 or out-of-band subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other than /64 are not supported.

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

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

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

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

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

Enable AAA Config – Select this check box to include AAA configurations from the Manageability tab as part of the device startup config post bootstrap.

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

Copy-paste the running-config to a **freeform config** field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. .

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

The format of the scope should be defined as:

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

For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

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

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

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

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

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

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

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

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

The backup configuration files are stored in the following path in Nexus Dashboard Fabric Controller: `/usr/local/cisco/dcn/dcnm/data/archive`

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



Note To trigger an immediate backup, do the following:

- 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.

11. Click on the **Fabric** to view summary in the slide-in pane. Click on the Launch icon to view the Fabric Overview.

Overview of Tenant Routed Multicast

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

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

A multicast-enabled data center fabric is typically part of an overall multicast network. Multicast sources, receivers, and multicast rendezvous points might reside inside the data center but also might be inside the

campus or externally reachable via the WAN. TRM allows a seamless integration with existing multicast networks. It can leverage multicast rendezvous points external to the fabric. Furthermore, TRM allows for tenant-aware external connectivity using Layer-3 physical interfaces or subinterfaces.

For more information, see the following:

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

Overview of Tenant Routed Multicast with VXLAN EVPN Multi-Site

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

The following two use cases are supported:

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

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

Tenant Routed Multicast with VXLAN EVPN Multi-Site Operations

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

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

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

Configuring TRM for Single Site Using Cisco Nexus Dashboard Fabric Controller

This section 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 DCNM

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.

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).

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. You can change from config-profile mode to CLI mode and from CLI mode to config-profile mode before you apply the overlay configurations on switches in the fabric.



Note Overlay-mode CLI is available only for Easy and eBGP Vxlan Fabrics.

After you upgrade from releases earlier than Cisco Nexus Dashboard Fabric Controller Release 12.0.1a, the existing config-profile mode will function the same.

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

If the switch has CLI-based overlays, you can import it in **config-profile** or **cli** overlay mode. If you set the overlay mode as **config-profile** the CLI-based overlays are converted to config-profile overlays.

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.
- 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)	<ul style="list-style-type: none"> • int_trunk_host • int_port_channel_trunk_host • int_vpc_trunk_host
access (standalone, po, and vPC PO)	<ul style="list-style-type: none"> • int_access_host • int_port_channel_access_host • int_vpc_access_host
dot1q-tunnel	<ul style="list-style-type: none"> • int_dot1q_tunnel_host • int_port_channel_dot1q_tunnel_host • int_vpc_dot1q_tunnel_host
routed	int_routed_host
loopback	int_freeform
sub-interface	int_subif
FEX (ST, AA)	<ul style="list-style-type: none"> • int_port_channel_fex • 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.

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.

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 62](#).

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 62](#).

8. Click **Save**.

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

9. From the Fabric Overview **Actions** drop-down list, select **Recalculate Config**.

The **Deploy Config** option can also be used for deployment. However, the **Recalculate Config** option identifies mismatch between the intended and running configuration *across all* fabric switches.

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:

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](#).

- 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.

- Now, optionally, click **Preview** to look at the configuration that will be pushed to the switch.
- 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-XPB-128
- GCM-AES-XPB-256

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

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

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

Step 4 Click a fabric 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 85](#).

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	<p>Enter the IP address of the switch.</p> <p>You can import more than one switch by providing the IP address range. For example: 10.10.10.40-60</p> <p>The switches must be properly cabled and reachable to the Cisco Nexus Dashboard Fabric Controller server and the switch status must be manageable.</p>

Field	Description
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 126](#) and [Templates, on page 333](#).



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.

- 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.

- 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.

- 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 137](#).

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 137](#).



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.

- 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.

- 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 145](#).

External Fabrics

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

- An external fabric is a monitor-only or managed mode fabric. Nexus Dashboard Fabric Controller supports only the monitor mode for Cisco IOS-XR family devices.
Cisco Nexus Dashboard Fabric Controller allows you to add or delete switches, when the external fabric is in the monitor mode.
- 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\$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

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 the fabric name 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. Nexus Dashboard Fabric Controller supports only the monitor mode for Cisco IOS-XR family devices.

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

When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches. If you click **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.

- 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 *Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - MPLS SR and LDP Handoff* chapter.

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

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

Enable 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, routed, or SVI interfaces), in addition to management of switches with out-of-band connectivity (aka reachable over switch mgmt0 interface). The only requirement is that for Inband managed switches, there should be IP reachability from Nexus Dashboard Fabric Controller to the switches via the eth2 aka inband interface. For this purpose, static routes may be needed on the Nexus Dashboard Fabric Controller, that in turn can be configured via the Administration->Customization->Network Preferences option. After enabling Inband management, during discovery, provide the IPs of all the switches to be imported using Inband Management and set maximum hops to 0. Nexus Dashboard Fabric Controller has a pre-check that validates that the Inband managed switch IPs are reachable over the eth2 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 111](#).

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. Additionally, the PTP Source Loopback Id and PTP Domain Id fields are editable. For more information, see [Precision Time Protocol for External Fabrics and LAN Classic Fabrics, on page 109](#).

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

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

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

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

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

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

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

Step 6 Fill up the **Resources** tab as shown below.

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 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 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.

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

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

Note Cisco 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 should be between 8 and 30.

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

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

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

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

Copy-paste the running-config to a **freeform config** field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. For more information, see [Enabling Freeform Configurations on Fabric Switches](#), on page 59.

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 **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 the 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. The allowed roles are Edge Router and Core Router. Assign the Core Router role for a Multi-Site Inter-Fabric Connection (IFC) and the Edge Router role for a VRF Lite IFC between the external fabric and VXLAN fabric border devices.



Note Changing of switch role is allowed only before executing **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 85](#) 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.

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.

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

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)# aaa authentication login default local
switch (config)# aaa authorization exec default local none
switch (config)# username admin privilege secret <password>
```



```
switch (config)# aaa new-model
switch (config)# session-id-common
```

Step 2 Run the following command in Nexus Dashboard Fabric Controller console to perform an SNMP walk.

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

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

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

Configuring Arista Devices for Discovery

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/0smUYydKmfuIjgn/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 Cisco IOS-XR Devices for Discovery

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

```

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

```



Note SNMP password should be same as password for username.

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

Configuration and Verification of Cisco IOS-XR Devices

```
RP/0/RSP0/CPU0:ios(config)#snmp-server view view_name cisco included
RP/0/RSP0/CPU0:ios(config)#snmp-server view view_name mib-2 included
RP/0/RSP0/CPU0:ios(config)#snmp-server group group_name v3 auth read view_name write view_name
RP/0/RSP0/CPU0:ios(config)#snmp-server user user_name group_name v3 auth md5 password_priv
des56 password SystemOwner
RP/0/RSP0/CPU0:ios(config)#commit Day MMM DD HH:MM:SS Timezone
RP/0/RSP0/CPU0:ios(config)#
RP/0/RSP0/CPU0:ios(config)#show run snmp-server Day MMM DD HH:MM:SS Timezone snmp-server
user user_name group1 v3 auth md5 encrypted 10400B0F3A4640585851_priv des56 encrypted
000A11103B0A59555B74 SystemOwner
snmp-server view view_name cisco included
snmp-server view view_name mib-2 included
snmp-server group group_name v3 auth read view_name write view_name
```

Discovering Non-Nexus Devices in an External Fabric

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

Before you begin

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

Procedure

- Step 1** Click **Add switches** in the **Actions** pane.
- Step 2** Enter values for the following fields under the **Discover Existing Switches** tab:

Field	Description
Seed IP	<p>Enter the IP address of the switch.</p> <p>You can import more than one switch by providing the IP address range. For example: 10.10.10.40-60</p> <p>The switches must be properly cabled and connected to the Nexus Dashboard Fabric Controller server and the switch status must be manageable.</p>
Device Type	<ul style="list-style-type: none"> Choose IOS XE from the drop-down list for adding Cisco CSR 1000v, Cisco ASR 1000 Series routers, or Cisco Catalyst 9000 Series Switches. Choose IOS XR from the drop-down list for adding ASR 9000 Series Routers, Cisco NCS 5500 Series Routers, IOS XR Release 6.5.3 or Cisco 8000 Series Routers. <p>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.</p> <ul style="list-style-type: none"> Choose Other from the drop-down list for adding non-Cisco devices, like Arista switches.

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

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

Set the password of the device in the **LAN Credentials** window if the password is not set. To navigate to the **LAN Credentials** window from the Cisco 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**.

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 85.

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 chapter 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.
- IPFM Easy 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.

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 94](#). For more information about Easy_Fabric_IPFM template, see [Creating an IPFM Easy Fabric, on page 97](#).

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 221](#).



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 119](#).

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 58](#).

If you want to edit or delete an IPFM fabric, see [Editing an IPFM Fabric, on page 104](#) or [Deleting an IPFM Fabric, on page 104](#) respectively.

4. Edit the existing interfaces as required. For more information, see [Editing an Interface for IPFM Fabrics, on page 107](#). For more information about any new logical interfaces, see [Creating an Interface for IPFM Fabrics, on page 105](#).

Creating an IPFM Classic Fabric

This section describes the procedure to create an IPFM classic fabric from the **IPFM_Classic 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** 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.

Enable Performance Monitoring - Select this check box to monitor the performance of the fabric.

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 Classic 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 eth2, that is, the inband 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 eth2 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 111](#).

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 59.

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 119](#).

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 Easy 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.

Step 5

Click the **Multicast** tab. The fields in this tab are:

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)*.

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 102](#) 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 102](#) 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 56](#).

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 59.

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 119](#).

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.

For more information, see [LAN Fabrics, on page 33](#).

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**.

For more information, see [LAN Fabrics, on page 33](#).

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 Nexus Dashboard Fabric Controller 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 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 bdpuguard
 - 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 107](#).

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 107](#).

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 108](#)

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 244](#).

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 105](#).

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 108](#).

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**.
-

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#~products>.



Note PTP global configuration is supported with Cisco Nexus 3000 Series switches; however, PTP and ttag configurations are not supported.

For more information, see the *Configuring PTP* chapter in *Cisco Nexus 9000 Series NX-OS System Management Configuration Guide* and *Cisco 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 58](#).

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

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 inband switch using seed IP. If there are multiple VRFs for same seed IP, then no intent will be learnt for seed interfaces. You must create intent/configuration manually.

After configuring/editing the Fabric settings, you must **Deploy Config**. You cannot change the Inband Mgmt settings after you import inband managed switches to the Fabric. If you uncheck the checkbox, the following error message is generated.

```
Inband IP <<IP Address>> cannot be used to import the switch,  
please enable Inband Mgmt in fabric settings and retry.
```

After the switches are imported to the Fabric, you must manage the interfaces to create intent. Create the intent for the interfaces that you're importing the switch. Edit/update the Interface configuration. When you try to change the Interface IP, for this inband managed switch, an error message is generated:

```
Interface <<interface_name>> is used as seed or next-hop egress interface  
for switch import in inband mode.  
IP/Netmask Length/VRF changes are not allowed for this interface.
```

While managing the interfaces, for switches imported using inband management, you cannot change the seed IP for the switch. The following error will be generated:

```
<<switch-name>>: Mgmt0 IP Address (<ip-address>) cannot be changed,  
when is it used as seed IP to discover the switch.
```

Create a policy for next-hop interfaces. Routes to Nexus Dashboard Fabric Controller from 3rd party devices may contain multiple interfaces, known as ECMP routes. Find the next-hop interface and create an intent for the switch. Interface IP and VRF changes are not allowed.

If inband management is enabled, during Image management, eth2 IP address is used to copy images on the switch, in ISSU, EPLD, RPM & SMU installations flows.

If you imported the switches using inband connectivity in the fabric, and later disable the inband Mgmt in the Fabric settings after deployment, the following error message is generated:

```
The fabric <<fabric name>> was updated with below message:
Fabric Settings cannot be changed for Inband Mgmt, when switches are already imported
using inband Ip. Please remove the existing switches imported using Inband Ip from the
fabric,
then change the Fabric Settings.
```

However, the same fabric can contain switches imported using both inband and out-of-band connectivity.

Enhanced Role-based Access Control

From Cisco Nexus Dashboard Fabric Controller Release 12.0.1a, you can create security domains in Nexus Dashboard, and map them with the user roles for Cisco Nexus Dashboard Fabric Controller. A user can have different roles on different security domains.

Cisco Nexus Dashboard Fabric Controller Release 12.0.1a has the following roles and privileges:

Roles	Privileges
NDFC Access Admin	Read/Write
NDFC Device Upgrade Admin	Read/Write
NDFC Network Admin	Read/Write
NDFC Network Operator	Read
NDFC Network Stager	Read/Write



Note In any window, the actions that cannot be performed 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 only if you are a user with the **NDFC 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 external fabrics if it isn't associated with policy.

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.
- 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.

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 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 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.

When you upgrade from Cisco Nexus Dashboard Fabric Controller Release 11.x, each fabric is mapped to an auto-generated 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 you created in Cisco Nexus Dashboard Fabric Controller Release 11.x. However, if you are creating a new AV-pair, use the format mentioned above. Ensure that the shell:domains should not have any spaces.

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.

Golden Backup

You can now mark the backups that you don't want to delete even after you reach the archiving limit. These backups are the golden backups. You can't delete golden backups of fabrics. However, Cisco Nexus Dashboard Fabric Controller archives only up to 10 golden backups. You can mark a backup as golden backup while restoring the fabric. To mark a backup as golden backup, perform the following steps from the Web UI:

Procedure

Step 1 Choose a fabric and choose **Fabrics > Fabric Overview > More > Backup Fabric**.

The **Backup** tab appears.

Step 2 On main window, choose **Actions > Configure Backup**.

The **Scheduled Archive** window appears.

Step 3 Choose the time period from where you want to choose the backup.

Valid values are **1m**, **3m**, **6m**, **YTD**, **1y**, and **All**. You can zoom into the graph. By default the backup information for **1m**, which is one month, appears. You can also choose a custom date range. The backup information includes the following information:

- Backup date
- Total number of devices
- Number of devices in sync
- Number of devices out of sync

Step 4 Choose the backup you want to mark as golden by clicking the backup.

You can choose the automatic or manual backup. These backups are color-coded. Automatic backups are indicated in blue color. Manual backups are indicated in midnight blue color. Golden backups are indicated in orange color. The automatic backups have only the versions in their names. Whereas the manual backups have tag names, which you gave when you initiated a manual backup, along with the version in the backup name. Hover over a backup to see the name. The automatic backup is initiated from the **Backup** tab in the **Fabric Overview** window. The manual backup is initiated by clicking **Backup Now** from the **Actions** pane in the backup tab.

Step 5 Navigate to switch window, choose check box for required switch name, choose **Switch > Switch Overview > Backup > Actions > Mark as golden backup** to mark golden backup.

A confirmation dialog box appears.

Step 6 Click **Yes**.

Step 7 Continue with rest of the fabric restore procedure as mentioned in the *Restoring Fabrics* section or exit the window.

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 . Refer to <i>Golden Backup</i> section for more details.
Remove as golden	To remove existing backup from golden backup, choose Remove as golden , a confirmation window appears, click Confirm .

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.
- 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, and the loopback interfaces are also configured for the corresponding VRFs.
- For host-to-host OAM, ensure that the Networks with VLAN are configured along

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:

- From the **Source Host IP** drop-down list, choose the source switch.
- From the **Destination Host IP** drop-down list, choose the destination switch.
- From the **VRF** drop-down list, choose or enter the VRF details.
- Check the **All Path 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:

- In the **Source Switch** field, enter the IP address of the source host.
- In the **Destination Switch** field, enter the IP 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.

Fabric Overview

The **Actions** drop-down list at the Fabric level allows you to perform the following:

Actions	Description
Edit Fabric	<ul style="list-style-type: none"> • To edit a fabric, choose Actions > Edit Fabric. • The Edit fabric window appears, do necessary changes and click Save.
Add Switches	Refer to section Adding Switches to a Fabric for more information.
Recalculate Config	Refer to section Recalculating and Deploying Configurations for more information.
Preview Config	Refer to section Previewing Switches for more information.
Deploy Config	<ul style="list-style-type: none"> • To deploy configuration changes, choose Actions > Deploy Config. • A progress window appears and confirmation message is displayed.
More	
Deployment Enable	<ul style="list-style-type: none"> • From Fabrics Overview, choose Actions on main tab, choose More > Deployment Enable. • A confirmation window appears, click OK.
Deployment Disable	<ul style="list-style-type: none"> • From Fabrics Overview, choose Actions on main tab, choose More > Deployment Disable. • A confirmation window appears, click OK.
Backup Fabric	Refer to Backup Fabric section for more information.
Restore Fabric	Refer to Restoring Fabric section for more information.
VXLAN OAM	<p>Refer to VXLAN OAM, on page 117 section for more information.</p> <p>Note 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.</p>

Actions	Description
Configure End Point Locator	<p>The Endpoint Locator (EPL) feature allows real-time tracking of endpoints within a data center. For more information, see Endpoint Locator , on page 204.</p> <p>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.</p>

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 if 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 153](#).

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 164](#).

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.
- 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.

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. 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: <ul style="list-style-type: none"> • Creating Inter-Fabric Links, on page 129 • Creating Intra-Fabric Links, on page 127
Edit	Allows you to edit the selected fabric.
Delete	Allows you to delete the selected fabric.

Action Item	Description
Import	<p>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.</p> <p>Note</p> <ul style="list-style-type: none"> You cannot update existing links. The Import Links icon is disabled for external fabric.
Export	<p>Choose the link and select Export to export the links in a CSV file.</p> <p>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.</p>

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 state 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 be converted into a config, but will not be pushed into the switch. After Save & Deploy, it will reflect in the running configuration.

Disable BFD Echo on Source Interface and Disable BFD Echo on Destination Interface – Select the check box to disable BFD echo packets on source and destination interface.

Note that the BFD echo fields are applicable only when you have enabled BFD in the fabric settings.

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

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 explained:

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

Link Sub-Type – This field populates the IFC type. Choose VRF_LITE, MULTISITE_UNDERLAY, or MULTISITE_OVERLAY from the drop-down list.

The Multi-Site options are explained in the Multi-Site use case.

For information about VXLAN MPLS interconnection, see the *Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - MPLS SR and LDP Handoff* chapter.

For information about routed fabric interconnection, see the *Creating Inter-Fabric Links Between a Routed Fabric and an External Fabric* section in the *Configuring a Fabric with eBGP Underlay* chapter.

Link Template: The link template is populated.

The templates are autopopulated with corresponding pre-packaged default templates that are based on your selection.

Note You can add, edit, or delete user-defined templates. See [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 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 **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 via MSD fabric, IFCs are automatically created between the (VXLAN fabric) border/BGW device and connected (external fabric) edge router/core device. When you remove the ER/core/border/BGW device, the corresponding IFCs (link PTIs) to/from that switch are deleted on Nexus Dashboard Fabric Controller. Subsequently, 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 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:

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.

Field	Description
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 Adding a Policy
Edit Policy	<p>Choose a policy from the table and choose Edit Policy to modify the policy.</p> <p>Note</p> <ul style="list-style-type: none"> The policies in the italics font cannot be edited. The value under the Editable and Mark Deleted columns for these policies is false. A warning appears when you edit a policy whose Mark Deleted value is set to <i>true</i>. 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	<p>Choose policies from the table and choose Delete Policy to delete the policies.</p> <p>Note A warning appears when you delete policies whose Mark Deleted values are set to <i>true</i>.</p>
Generated Config	Choose policies from the table and choose Generated Config to view the delta of configuration changes made by every user.

Action Item	Description
Push Config	<p>Choose policies from the table and choose Push Config to push the policy configuration to the device.</p> <p>Note</p> <ul style="list-style-type: none"> • This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric. • A warning appears if you push configuration for a Python policy. • A warning appears when you push configurations for policies whose Mark Deleted values are set to <i>true</i>.

Event Analytics

Event Analytics includes the following topics:

- [Alarms, on page 295](#)
- [Events, on page 300](#)
- [Accounting, on page 304](#)

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 57](#) 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 135 .
Edit	Allows you to edit the selected VRF. 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. To import VRF information, choose Import . Browse the directory and select the <code>.csv</code> file that contains the VRF information. Click Open . The VRF information is imported and displayed in the VRFs tab of the Fabric Overview window.

Action Item	Description
Export	<p>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.</p> <p>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.</p> <p>Note You can use the exported .csv file for reference or use it as a template for creating new VRFs.</p>
Delete	<p>Allows you to delete a selected VRF.</p> <p>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.</p>

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 a 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** Click **Actions** and choose **Create**.
The **Create VRF** window appears.
- Step 2** Enter the required details in the mandatory fields. The available fields vary slightly 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 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.

VRF Extension Template - A universal extension template is autopopulated. This allows you to extend this network to another fabric. The methods are VRF Lite, Multi Site, and so on. The template is applicable for border leaf switches and BGWs.

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

a) The fields on the **General** tab are:

VRF Vlan Name - Enter the VLAN name for the VRF.

VRF Description - Enter a description for the VRF.

VRF Intf Description - Enter a description for the VRF interface.

b) Click the **Advanced** tab to optionally specify the advanced profile settings. The fields on this tab are autopopulated. The fields on the **Advanced** tab are:

VRF Intf 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 too.

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.

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

RP Address – Specifies the IP address of the RP.

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

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

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

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

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

Enable TRM BGW MSite - Select the check box to enable TRM on Border Gateway Multisite.

Advertise Host Routes - Enable 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.

To allow inter-subnet communication between end hosts in different VXLAN fabrics, where the subnets are present in both fabrics, you must disable the **Advertise Default Route** feature (clear the **Advertise**

Default Route checkbox) for the associated VRF. This will result in /32 routes for hosts being seen in both fabrics. For example, Host1 (VNI 30000, VRF 50001) in Fabric1 can send traffic to Host2 (VNI 30001, VRF 50001) in Fabric2 only if the host route is present in both fabrics. When a subnet is present in only one fabric then default route is sufficient for inter-subnet communication.

Config Static 0/0 Route - Enable 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 - Select the encryption type from this drop-down list.

Step 3 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 switch name.
Status	Specifies the status of the 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.

Field	Description
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. Note 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	<p>Allows you to view the deployment and policy change history of the selected VRF.</p> <p>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.</p> <p>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.</p> <p>To view the history of a VRF attachment, select the check box next to the VRF name and choose the History action. 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.</p>
Edit	<p>Allows you to view or edit the VRF attachment parameters such as interfaces that you want to attach to the selected VRF.</p> <p>To edit the VRF attachment information, select the check box next to the VRF name that you want to edit and choose the Edit action. 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.</p>

Action Item	Description
Preview	<p>Allows you to preview the configuration of the VRF attachments for the selected VRF.</p> <p>Note This action is not allowed for attachments that are in deployed or NA status.</p> <p>To preview the VRF, select the check box next to the VRF name and choose the Preview action. The Preview Configuration window for the fabric appears.</p> <p>You can preview the VRF attachment details such as the VRF name; fabric name; switch name, serial number, IP address, and role; the VRF status, pending configuration, and progress of the configuration. You can also click the lines link in the Pending Config column to view the lines for which the configuration is pending. Click Close.</p>
Deploy	<p>Allows you to deploy the pending configuration of the VRF attachments, for example, interfaces, for the selected VRF.</p> <p>Note This action is not allowed for attachments that are in deployed or NA status.</p> <p>To deploy a VRF, select the check box next to the VRF name and choose the Deploy action. The Deploy Configuration window for the fabric appears.</p> <p>You can view the details such as the VRF name; fabric name; switch name, serial number, IP address, and role; the VRF status, pending configuration, and progress of the configuration. You can also 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.</p>
Import	<p>Allows you to import information about VRF attachments for the selected fabric.</p> <p>To import the VRF attachments information, choose Import. Browse the directory and select the <code>.csv</code> 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.</p>

Action Item	Description
Export	<p>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.</p> <p>To export VRF attachments information, choose the Export action. 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.</p>
Quick Attach	<p>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.</p> <p>To quickly attach any attachment to a VRF, choose the Quick Attach action. A message appears to inform you that the attach action was successful.</p>
Quick Detach	<p>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.</p> <p>To attach any attachment to a VRF quickly, choose the Quick Detach action. A message appears to inform you that the detach action was successful.</p>

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 133](#).

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 57](#) section.

For information about creating interface groups and attaching networks, see [Interface Groups, on page 265](#).

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 Networks for the Standalone Fabric, on page 143 .
Edit	Allows you to view or edit the selected network parameters. 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. To import network information, choose Import . Browse the directory and select the <code>.csv</code> 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 <code>.csv</code> 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. 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. Note You can use the exported <code>.csv</code> file for reference or use it as a template for creating new networks. Before importing the file, update new records in the <code>.csv</code> 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.

Action Item	Description
Delete	<p>Allows you to delete the network.</p> <p>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.</p>
Add to interface group	<p>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.</p> <p>To add the selected networks to the interface group that you want, click Add to interface group action.</p> <p>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.</p> <p>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.</p> <p>In the Add to interface group window, click Save to save the changes and close the window or click Cancel to discard the changes.</p> <p>The interface group is displayed in a column in the Networks tab of the Fabric Overview window.</p>
Remove from interface group	<p>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.</p> <p>To remove the selected networks to the interface group that you want, click Remove from interface group action.</p> <p>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.</p> <p>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.</p> <p>The interface group are removed from the column in the Networks tab of the Fabric Overview window.</p>

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.

Field	Description
VRF Name	Specifies the name of the Virtual Routing and Forwarding (VRF).
IPv4 Gateway/Suffix	Specifies the IPv4 address with subnet.
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 Networks for the Standalone Fabric

To create 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, you do not require a VRF. For more information about VRFs, see [VRFs, on page 133](#).

Procedure

-
- Step 1** Click **Actions** and choose **Create**.
- The **Create Network** window appears.
- Step 2** Enter the required details in the mandatory fields. The available fields vary slightly based on the fabric type. 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 (-). 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).
- 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 methods are VRF Lite, Multi Site, and so on. The template is applicable for border leaf switches and BGWs.

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.

a) 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.

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.

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.

b) 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 only 128. If all networks are deployed on all switches, you need not use a different multicast group per L2 VNI or a network. Therefore, multicast group for all networks in a fabric remains same. If a new multicast group address is required, you can generate it by clicking the **Generate Multicast IP** button.

DHCPv4 Server 1: Enter the DHCP relay IP address of the first DHCP server.

DHCPv4 Server 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: Select the check box to enable TRM.

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

L2 VNI Route Target Both Enable: Select the check box to enable automatic importing and exporting of route targets for all L2 virtual networks.

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

Step 3 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.

Network 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 > Networks > Network Attachments**.
- Choose **LAN > Fabrics**. Double-click on a fabric to open **Fabric Overview > Networks > Network Attachments**.

Use this window to attach attachments such as fabrics and interfaces to a network.

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 7: Network Attachments Actions and Description

Action Item	Description
History	<p>Allows you to view the deployment and policy change history of the selected network.</p> <p>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.</p> <p>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.</p> <p>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. 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.</p>
Edit	<p>Allows you to view or edit the network attachment parameters such as interfaces that you want to attach to the selected network.</p> <p>To edit the network attachment information, select 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.</p>
Preview	<p>Allows you to preview the configuration of the network attachments for the selected network.</p> <p>Note This action is not allowed for attachments that are in deployed or NA status.</p> <p>To preview the network, select the check box next to the network name and choose the Preview action. The Preview Configuration window for the fabric appears.</p> <p>You can preview the network attachment details such as the network name; fabric name; switch name, serial number, IP address, and role; the network status, pending configuration, and progress of the configuration. You can also click the lines link in the Pending Config column to view the lines for which the configuration is pending. Click Close.</p>

Action Item	Description
Deploy	<p>Allows you to deploy the pending configuration of the network attachments, for example, interfaces, for the selected network.</p> <p>Note This action is not allowed for attachments that are in deployed or NA status.</p> <p>To deploy a network, select the check box next to the network name and choose the Deploy action. The Deploy Configuration window for the fabric appears.</p> <p>You can view the details such as the network name; fabric name; switch name, serial number, IP address, and role; the network status, pending configuration, and progress of the configuration. You can also 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.</p>
Import	<p>Allows you to import information about network attachments for the selected fabric.</p> <p>To import the network attachments information, choose Import. Browse the directory and select the <code>.csv</code> 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.</p>
Export	<p>Allows you to export the information about network attachments to a <code>.csv</code> 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.</p> <p>To export network attachments information, choose the Export action. 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.</p>

Action Item	Description
Quick Attach	<p>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.</p> <p>Note Interfaces cannot be attached to a network using this action.</p> <p>To quickly attach any attachment to a network, choose the Quick Attach action. A message appears to inform you that the attach action was successful.</p>
Quick Detach	<p>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.</p> <p>To quickly attach any attachment to a network, choose the Quick Detach action. A message appears to inform you that the detach action was successful.</p>

Table 8: 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 switch name.
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.

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 121](#).

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: <ul style="list-style-type: none"> • Sender • External Sender • Dynamic Receiver • External Receiver • 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 160 .

Action Item	Description
Edit Host Policy	<p>Allows you to view or edit the selected host policy parameters.</p> <p>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.</p> <p>Note 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.</p>
Delete Host Policy	<p>Allows you to delete user-defined host policies.</p> <p>Note</p> <ul style="list-style-type: none"> • Undeploy policies from all the switches before deleting them from Nexus Dashboard Fabric Controller. • Default policy can be undeployed from the switches on which it is deployed. However, Custom policy can be deleted and undeployed. • When you undeploy the default policies, all default policies are reset to have default permission (Allow). <p>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.</p> <p>A delete host policy successful message appears at the bottom of the page.</p>
Purge	<p>Allows you to delete all custom policies without selecting any policy check box.</p> <p>Note</p> <ul style="list-style-type: none"> • Undeploy policies from all switches before deleting them from Nexus Dashboard Fabric Controller. • 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	<p>Allows you to import host policies from a CSV file to Nexus Dashboard Fabric Controller.</p> <p>Note After import, all policies imported from a CSV file are applied to all managed switches automatically.</p> <p>To import a host policies, choose Import. Browse the directory and select the <code>.csv</code> file that contains the host policy configuration information. The policy will not be imported if the format in the <code>.csv</code> file is incorrect. Click Open. The imported policies are automatically deployed to all the switches in the fabric.</p>
Export	<p>Allows you to export host policies from Nexus Dashboard Fabric Controller to a <code>.csv</code> file.</p> <p>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 <code>.csv</code>.</p>
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	<p>Select this option to undeploy the selected policies.</p> <p>Select one or more check boxes next to the policy name. Select this option from the drop-down list to undeploy the selected policies.</p>
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	<p>The deployment of policies may fail due to various reasons. Select this option to deploy or undeploy all failed policies.</p> <p>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.</p>

Action Item	Description
Deployment History	<p>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.</p> <p>The policy name is shown in the Policy Name field. From the drop-down list, choose the switch on which this policy was deployed.</p> <p>The Deployment History pane displays the following fields.</p> <ul style="list-style-type: none"> • Policy Name - Specifies the selected policy name. • VRF - Specifies the VRF for the selected policy. • Switch Name - Specifies the name of the switch that the policy was deployed to. • 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 159. • 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. • Deployment Date/Time - Specifies the date and time at which the host policy was last updated. The format is <i>Day MMM DD YYYY HH:MM:SS Timezone</i>. • 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: <ul style="list-style-type: none"> • Sender • Receiver • Receiver-External • Receiver-Local
Operation	Specifies if the operation of the host policy. The policy has the following operations: <ul style="list-style-type: none"> • Permit • 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. <ul style="list-style-type: none"> • Create - The policy is deployed on the switch. • 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. The format is <i>Day MMM DD YYYY HH:MM:SS Timezone</i> .

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 162 .
Edit Host Alias	Allows you to view or edit the selected host alias parameters. 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. 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. To import host aliases, choose Import . Browse the directory and select the <code>.csv</code> 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. 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 <code>.csv</code> .

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 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 > 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. The host device role is either one of the following: <ul style="list-style-type: none"> • PIM • Sender • 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. 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 121](#).

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 If you want 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, then click the **active** hyperlink in **Flow Link State** for a particular multicast IP.

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 193](#).

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 17: NAT Field and Description

Field	Description
NAT	Specifies the NAT mode, that is, Ingress, Egress, or Ingress and Egress. For the Ingress NAT type, the following information is displayed: Ingress (S) – Specifies that ingress NAT is performed on the Sender Switch, also known as First Hop Router (FHR). Ingress (R) - Specifies that ingress NAT is performed on the Receiver Switch (also known as Last Hop Router (LHR)). 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 18: 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.
Multicast IP	Specifies the multicast IP address for the flow. Note 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. Click the active link to view the network diagram or topology of the Sender and Receiver. 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. 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 19: 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. Note 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	<p>Specifies reason for the inactive flow.</p> <p>Cisco Nexus Dashboard Fabric Controller determines the inactive flow if both the sender and receiver mroute exists with any of the following combinations.</p> <ul style="list-style-type: none"> • Receiver IIF is null • Receiver OIF is null • Sender IIF is null • Sender OIF is null <p>In this scenario, the switch will not have any fault reason. Therefore, there is no fault reason for such inactive flows.</p>
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 20: 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	<p>Specifies the flow link state, if it's allow or deny.</p> <p>Click the senderonly link to view the network diagram or topology of the Sender and Receiver.</p> <p>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.</p> <p>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.</p>
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.

Field	Description
Common Fields for IPFM and Generic Multicast Modes	
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 21: 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.
Flow Alias	Specifies the name of the Flow Alias.
Flow Link State	<p>Specifies the flow link state, if it's allow or deny.</p> <p>Click the receiveronly link to view the network diagram or topology of the Sender and Receiver.</p> <p>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.</p> <p>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.</p>
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.

Field	Description
Common Fields for IPFM and Generic Multicast Modes	
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 22: 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. <ul style="list-style-type: none"> • Create - The policy is deployed on the switch. • 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. <p>Note In adding or editing a flow policy, the default policer state is Enabled.</p>

Field	Description
Last Updated	Specifies the date and time at which the flow policy was last updated. The format is <i>Day MMM DD YYYY HH:MM:SS Timezone</i> .

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 23: 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 175 .
Edit Flow Policy	Allows you to view or edit the selected flow policy parameters. Note 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. 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. 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.

Field	Description
Delete Flow Policy	<p>Allows you to delete the user-defined flow policy.</p> <p>Note</p> <ul style="list-style-type: none"> • You cannot delete the default flow policies. • Undeploy policies from all switches before deleting them from Nexus Dashboard Fabric Controller. • You can select more than one flow policy to delete. <p>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.</p>
Purge	<p>Allows you to delete all the flow policies at a single instance.</p> <p>Note Undeploy policies from all switches before deleting them from Nexus Dashboard Fabric Controller.</p> <p>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.</p>
Import	<p>Allows you to import flow policies from a csv file.</p> <p>Note 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.</p> <p>After import, all policies imported from a csv file are applied to all managed switches automatically.</p> <p>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.</p>
Export	<p>Allows you to export flow policies to a csv file.</p> <p>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.</p>
Deploy Selected Policies	<p>Select this option to deploy only the selected policies to the devices. You can deploy other policies when required.</p> <p>Select one or more check boxes next to the policy name. Select this option to deploy selected policies to the switch.</p>

Field	Description
Deploy All Custom Policies	Select this option to deploy all the custom or user-defined policies at a single instance. 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. 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.
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. 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. The Deployment History pane displays the following fields: <ul style="list-style-type: none"> • Policy Name - Specifies the selected policy name. • VRF - Specifies the VRF for the selected policy. • Switch Name - Specifies the name of the switch that the policy was deployed to. • 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 174. • Action - Specifies the action that is performed on the switch for that flow policy. <ul style="list-style-type: none"> • Create - Implies that the policy has been deployed on the switch. • Delete - Implies that the policy has been undeployed from the switch. • 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 . • 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 24: 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.
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 25: 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.
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 26: 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 178 .
Edit Flow Alias	Allows you to view or edit the selected flow alias parameters. 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. 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. To import flow aliases, choose Import . Browse the directory and select the <code>.CSV</code> 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. 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 <code>.CSV</code> .

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 27: 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 179 .

Field	Description
Delete Static Flow	Allows you to delete the static flow. Select a static flow that you need to delete and click the Delete Static Flow action to delete the selected static flow.

Table 28: 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. 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 average Rx value.
Temperature Module	Specifies the peak Rx value.
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	Specifies the speed of switch.
Status	Specifies the status of switch.
Speed	Specifies the speed of the interface.
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	

Fields	Descriptions
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.

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 loopback 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. 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	<p>Select a mode from the table and choose Deployment History to view the deployment history of the selected mode.</p> <p>Deployment History shows the following fields.</p> <ul style="list-style-type: none"> • Switch Name—Specifies the name of the switch that the mode was deployed to. • VRF—Specifies the name of the VRF that mode was deployed to. • Group—Specifies the multicast group of the NAT mode. • Mode—Specifies the NAT mode, that is, ingress or egress. • Deployment Status—Displays the status of deployment. It shows if the deployment was Success or Failed. • 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 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. • 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. 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. 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	<p>Select a mode from the table and choose Deployment History to view the deployment history of the selected egress interface mapping.</p> <p>Deployment History shows the following fields.</p> <ul style="list-style-type: none"> • Switch Name—Specifies the name of the switch that the mode was deployed to. • Map Interface—Specifies the map interface for the egress interface mappings. • Max Replications—Specifies the maximum replications for the egress interface mappings. • Egress Interfaces—Specifies the name of the egress interface that the mapping is deployed to. • Deployment Status—Displays the status of deployment. It shows if the deployment was Success or Failed. • 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. • 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. • 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. 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	<p>Select a rule from the table and choose Deployment History to view the deployment history of the selected rule.</p> <p>Deployment History shows the following fields.</p> <ul style="list-style-type: none"> • Switch Name—Specifies the name of the switch that the rule was deployed to. • VRF—Specifies the VRF that the mapping belongs to. • Deployment Status—Displays the status of deployment. It shows if the deployment was Success or Failed. • 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. • 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. • 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.

Status Of: 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
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.
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](#).

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 29: 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 %	<p>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.</p> <p>You can configure the server to allot a dedicated percentage of bandwidth to unicast traffic. The remaining percentage is automatically reserved for multicast traffic.</p> <p>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 199.</p> <p>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 199.</p>

Field	Description
Reserve Bandwidth to Receiver Only	<p>Bandwidth reservation status specifies if the bandwidth deployment was success, or failed or not deployed.</p> <p>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.</p> <p>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 199.</p> <p>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 199.</p>
ASM/MASK	<p>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.</p> <p>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.</p> <p>The IP address and subnet mask in the ASM/MASK field define the multicast source.</p> <p>The ASM range is configured by specifying the IP address and the subnet mask.</p> <p>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 199.</p> <p>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 199.</p>

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 30: Switch Global Config Actions and Description

Action Item	Description
Edit NBM VRF Config	Allows you to edit the NBM VRF configuration. 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 31: 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 32: 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**.

Use the **IPFM VRF** window to create, edit, delete, and redeploy VRFs. You can view the deployment status and history of each VRF.

Table 33: IPFM VRF Table Fields and Description

Field	Description
Name	Specifies the name 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 . Click the Failed status to view more information about the Deployment Status, on page 199 .

Field	Description
Deployment History	Specifies the deployment history of the VRF. For default VRFs, the deployment history is displayed as Not Applicable . 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 34: IPFM VRF Actions and Description

Action Item	Description
Create VRF	Allows you to create a new VRF. 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, and click Save & Deploy to retain the changes and deploy or click Cancel to discard the changes. Note When you create a custom or non-default 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. For more information about deploying the policies manually, see Host Policies , on page 154 and Flow Policies .
Edit VRF	Allows you to edit a selected VRF. 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. 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. 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 35: 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.
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 36: 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 37: 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 38: VRF Actions and Description

Action Item	Description
Create VRF	Allows you to create a new VRF. 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. 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.

Action Item	Description
Delete VRF	Allows you to delete a selected VRF. 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.

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.



Important

- 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 needs to 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 eth2 interface.

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. This is supported only in clustered mode.

- EPL is supported on Multi-Site Domain (MSD).
- IPv6 underlay is supported.
- Support for high availability
- Support for endpoint data that is stored for up to 180 days, amounting to a maximum of 100 GB storage space.
- Support for optional flush of the endpoint data in order to start afresh.
- Supported scale: 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.

For more information about EPL, refer to the following sections:

Configuring Endpoint Locator

The Nexus Dashboard Fabric Controller OVA or the ISO installation comes with three interfaces:

- eth0 interface for external access
- eth1 interface for fabric management (Out-of-band or OOB)
- eth2 interface for in-band network connectivity

The eth1 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. For this purpose, the eth2 interface can be configured using the **appmgr update network-properties** command. Optionally, you can configure the eth2 interface during the Cisco Nexus Dashboard Fabric Controller installation.

If you need to modify the already configured in-band network (eth2 interface), run the **appmgr update network-properties** command again.



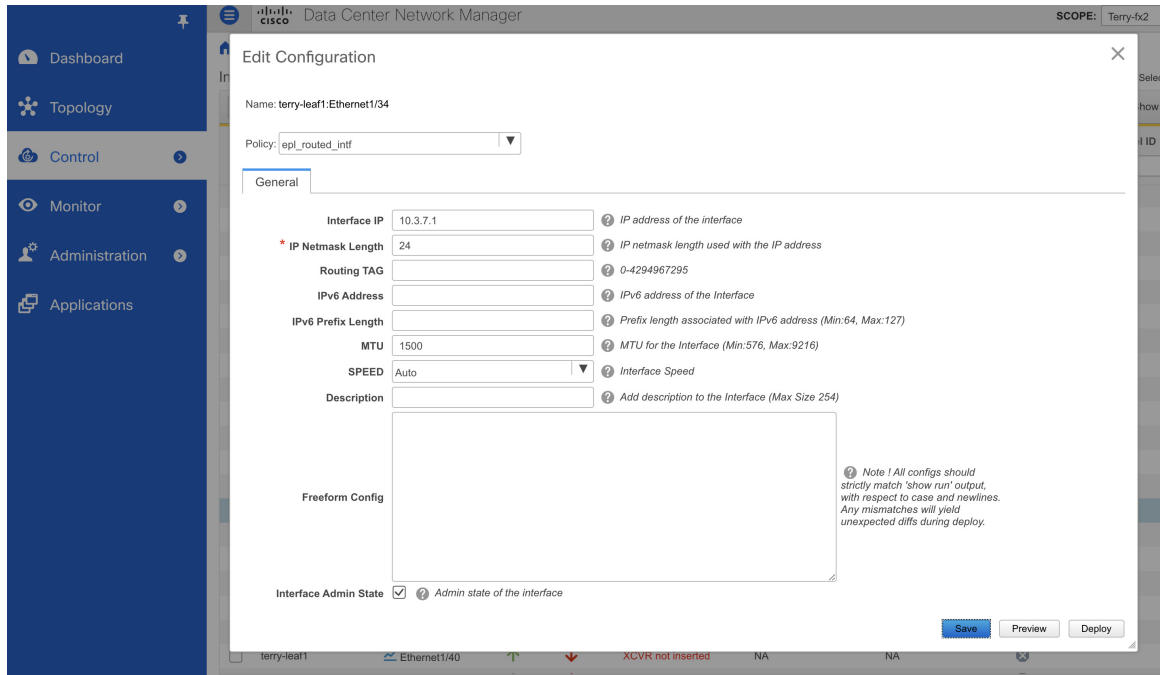
Note The setup of eth2 interface on the Nexus Dashboard Fabric Controller is a prerequisite of any application that requires the in-band connectivity to the devices within fabric. This includes EPL and Network Insights Resources (NIR).




Note For configuring EPL in standalone mode, you must add a single neighbor to EPL. Nexus Dashboard Fabric Controller eth2 IP address is EPL IP.

On the fabric side, for a standalone Nexus Dashboard Fabric Controller deployment, if the Nexus Dashboard Fabric Controller eth2 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 eth2 interface on the Nexus Dashboard Fabric Controller. The following image depicts an example scenario configuration:



In this example, the server with the Nexus Dashboard Fabric Controller VM is dual-attached to a vPC pair of switches that are named Site2-Leaf2 and Site2-Leaf3 respectively. VLAN 596 associated with the IP subnet 10.3.7.0/24 is employed for in-band connectivity. You can configure the vPC host port toward the server using the **interface vpc trunk host** policy as shown in the following image:

For the HSRP configuration on Site2-Leaf2, the **switch_freemom** policy may be employed as shown in the following image:

You can deploy a similar configuration on Site2-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 eth2 interface with the default gateway set to 10.3.7.1.

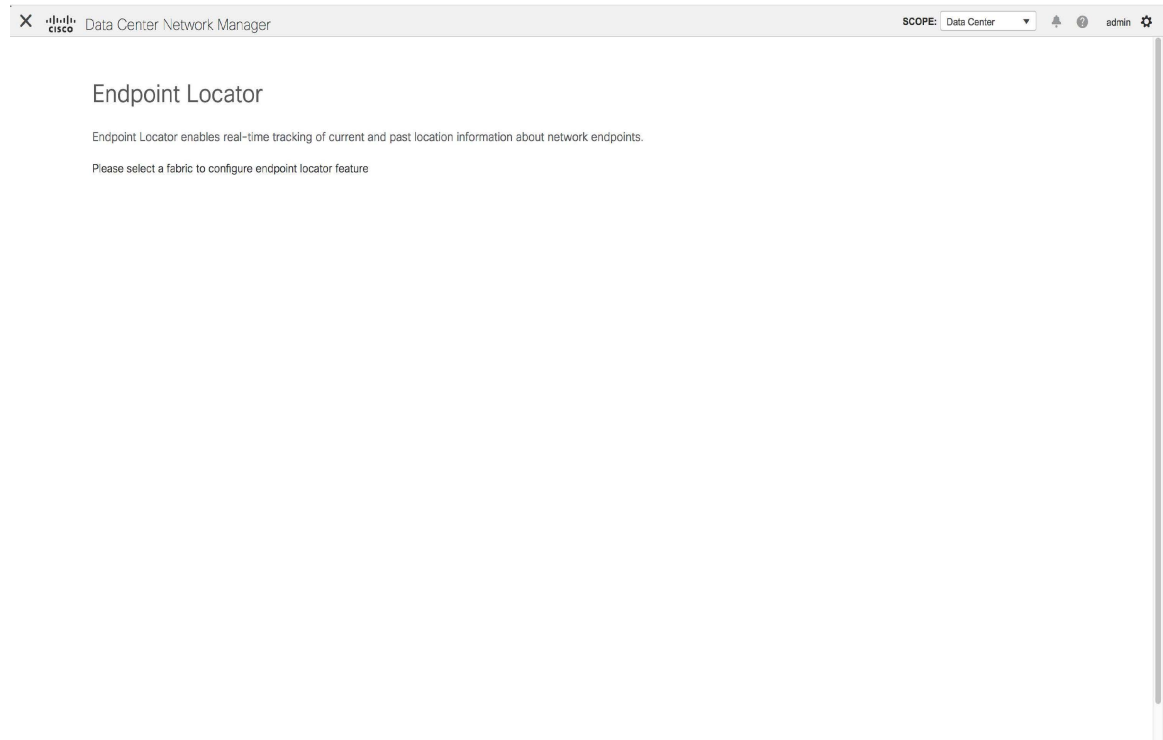
After you establish the in-band connectivity between the physical or virtual Nexus Dashboard Fabric Controller and the fabric, you can establish BGP peering.

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 eth2 gateway.



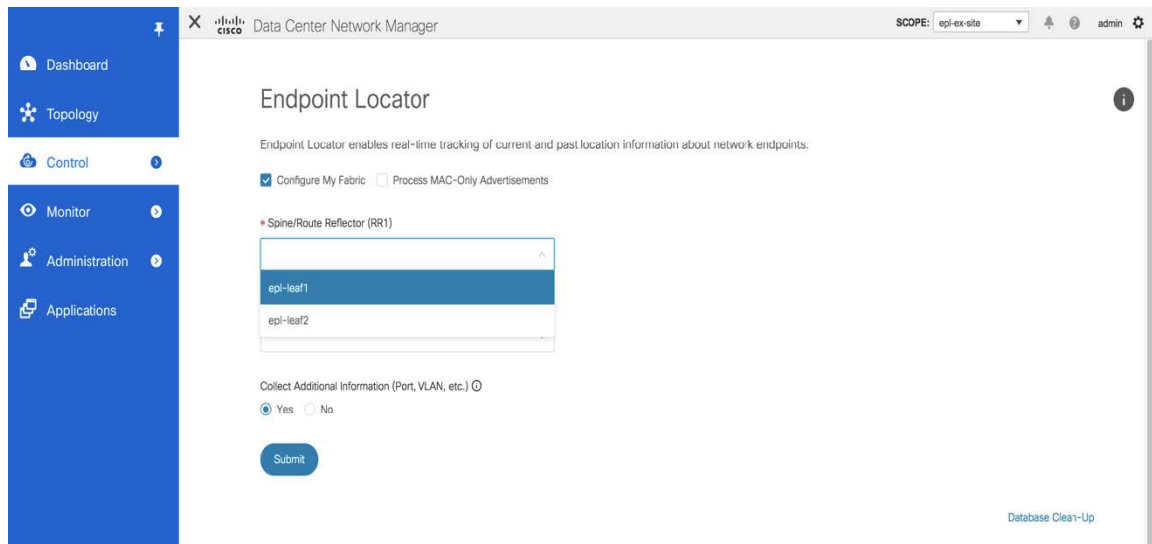
Note Cisco Nexus Dashboard Fabric Controller queries the BGP RR to glean information for establishment of the peering, like ASN, RR, IP, and so on.

To configure Endpoint Locator from the Cisco Nexus Dashboard Fabric Controller Web UI, choose **Endpoint Locator > Configure**. The **Endpoint Locator** window appears.



Select a fabric on which the endpoint locator feature should be enabled to track endpoint activity. 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 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 neighborhood, 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.

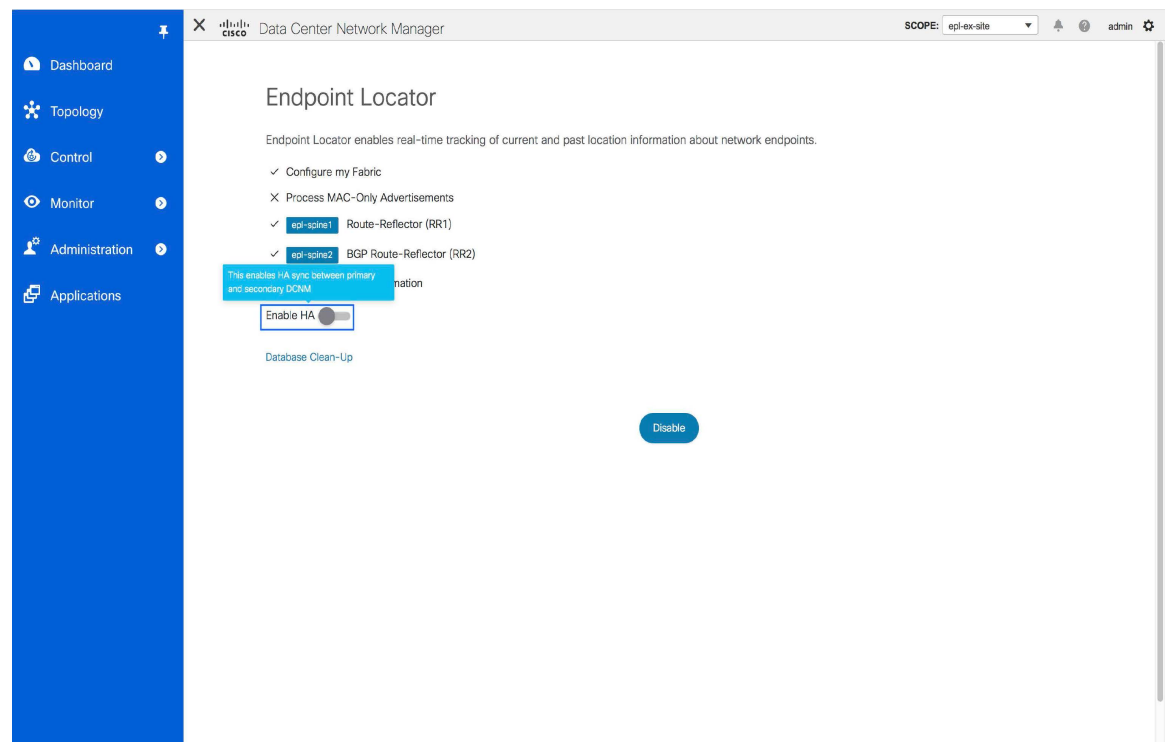
Once the appropriate selections are made and various inputs have been reviewed, click **Submit** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled.

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. For the native HA Nexus Dashboard Fabric Controller deployment, both the primary and secondary Nexus Dashboard Fabric Controller eth2 interface IPs will be added as BGP neighbors but only one of them will be active at any given time. 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 7](#).

Enabling High Availability

Consider a scenario in which EPL is enabled on a Nexus Dashboard Fabric Controller deployment that is in non-HA mode and then, Nexus Dashboard Fabric Controller is moved to HA-mode. In such scenarios, the **Enable HA** toggle appears on the **Endpoint Locator** window. Toggle the **Enable HA** knob to enable high availability sync between primary and secondary Nexus Dashboard Fabric Controller.



To enable high availability sync from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

Procedure

-
- Step 1** Choose **Endpoint Locator > Configure**.
- Step 2** Toggle the **Enable HA** button.
-

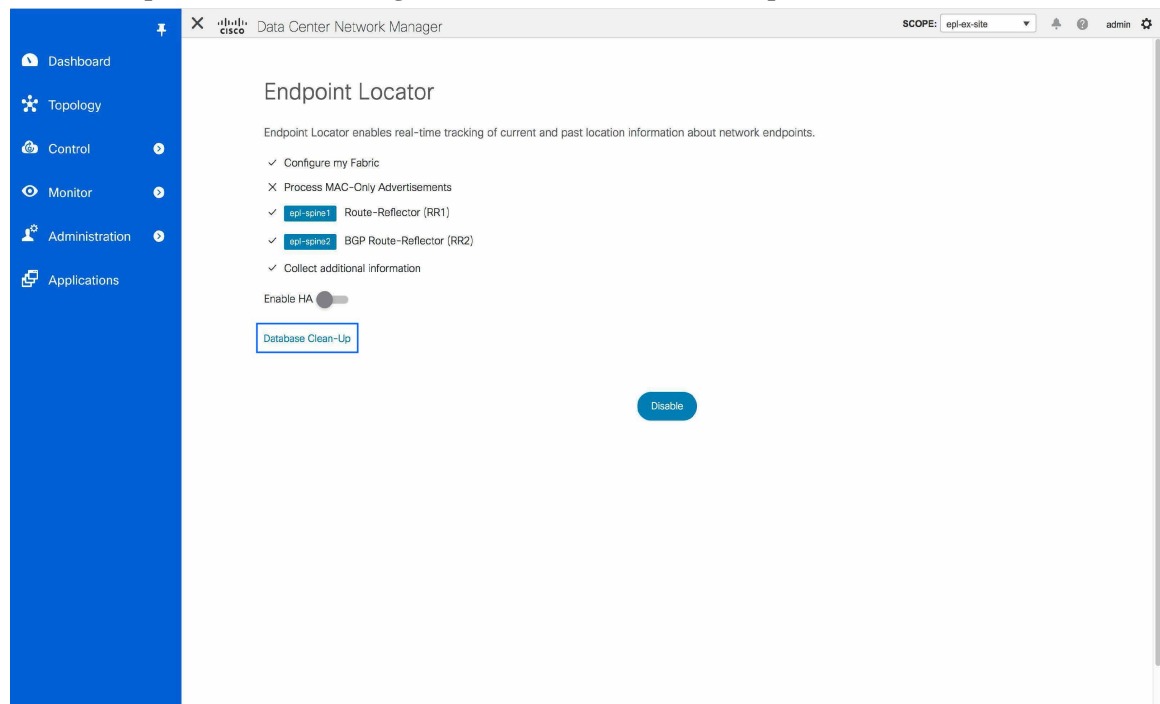
Flushing the Endpoint Database

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**.



A warning is displayed with a message indicating that all the endpoint information that is stored in the database will be flushed.

Step 2 Click **Delete** to continue or **Cancel** to abort.

Configuring Endpoint Locator in NDFC High Availability Mode



Note For configuring EPL in native HA mode, you must add 2 neighbors to EPL. EPL IP being Nexus Dashboard Fabric Controller Primary eth2 and Nexus Dashboard Fabric Controller Secondary eth2 address respectively.

For production deployments, a native HA pair of Nexus Dashboard Fabric Controller nodes is recommended. Since the Nexus Dashboard Fabric Controller active and standby nodes need to be Layer-2 adjacent, their respective eth2 interfaces should be part of the same IP subnet or vlan. In addition, both Nexus Dashboard Fabric Controller nodes should be configured with the same eth2 gateway. The recommended option is to connect the Nexus Dashboard Fabric Controller active and standby nodes to a vPC pair of nexus switches (they may be leafs) so that there is enough fault-tolerance in case of single link failure, single device or a single Nexus Dashboard Fabric Controller node failure.

The following example shows a sample output for the **appmgr update network-properties** command for a Cisco Nexus Dashboard Fabric Controller Native HA Appliance. In this example, 1.1.1.2 is the primary eth2 interface IP address, 1.1.1.3 is the standby eth2 interface IP address, 1.1.1.1 is the default gateway and 1.1.1.4 is the virtual IP (VIP) for inband.

On Cisco Nexus Dashboard Fabric Controller Primary appliance:

```
appmgr update network-properties session start
appmgr update network-properties set ipv4 eth2 1.1.1.2 255.255.255.0 1.1.1.1
appmgr update network-properties set ipv4 peer2 1.1.1.3
appmgr update network-properties set ipv4 vip2 1.1.1.4 255.255.255.0
appmgr update network-properties session apply
appmgr update ssh-peer-trust
```

On Cisco Nexus Dashboard Fabric Controller Secondary appliance:

```
appmgr update network-properties session start
appmgr update network-properties set ipv4 eth2 1.1.1.3 255.255.255.0 1.1.1.1
appmgr update network-properties set ipv4 peer2 1.1.1.2
appmgr update network-properties set ipv4 vip2 1.1.1.4 255.255.255.0
appmgr update network-properties session apply
appmgr update ssh-peer-trust
```

After the in-band connectivity is established from both the Primary and Secondary nodes to the Fabric, to configure endpoint locator in Nexus Dashboard Fabric Controller HA mode from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

Procedure

-
- Step 1** Choose **Control > Endpoint Locator > Configure**.
 - The **Endpoint Locator** window appears and the fabric configuration details are displayed.
 - Step 2** Select a fabric from the **SCOPE** dropdown list to configure endpoint locator in Nexus Dashboard Fabric Controller HA mode.
 - Step 3** Select the Route-Reflectors (RRs) from the drop-down lists.
 - Step 4** Select **Yes** under **Collect Additional Information** to enable collection of additional information such as PORT, VLAN, VRF etc. while enabling the EPL feature. If the No option is selected, this information will not be collected and reported by EPL.
 - Step 5** Click **Submit**.
-

What to do next

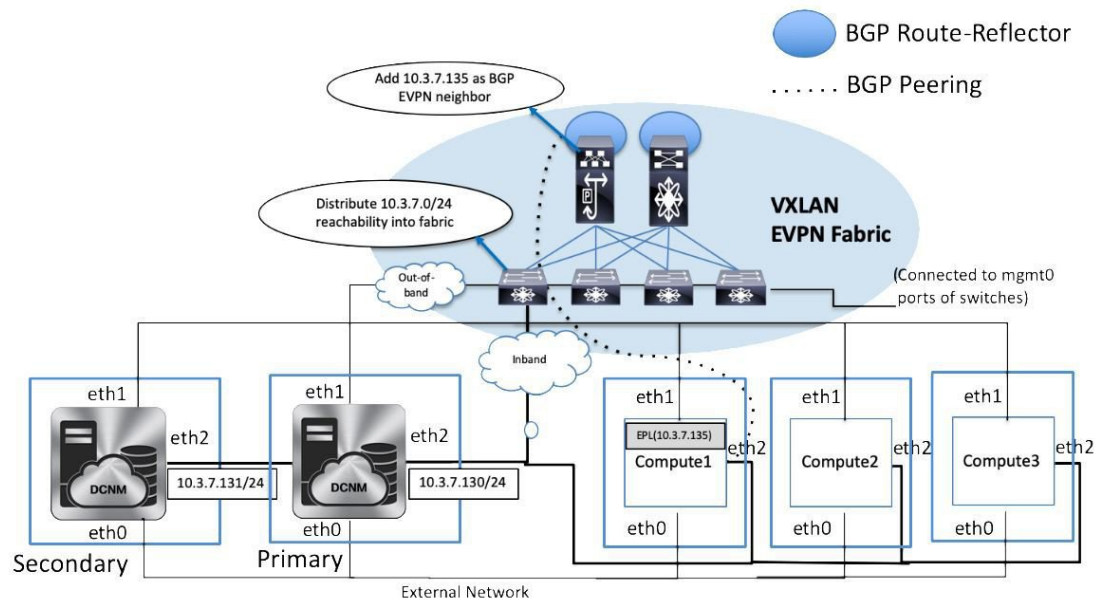
After you configure the Endpoint Locator in HA mode, you can view details such as Endpoint Activity and Endpoint History in the Endpoint Locator dashboard. To view these details, navigate to **Monitor > Endpoint Locator > Explore**.

Configuring Endpoint Locator in NDFC Cluster Mode



Note For configuring EPL in cluster mode, you must add a single neighbor to EPL. Nexus Dashboard Fabric Controller EPL container Inband IP address is EPL IP.

With the Nexus Dashboard Fabric Controller cluster mode deployment, in addition to the Nexus Dashboard Fabric Controller nodes, an additional 3 compute nodes are present in the deployment. For information about deploying applications in cluster mode, see *Cisco NDFC in Clustered Mode*.



In Nexus Dashboard Fabric Controller Cluster mode, all applications including EPL run on the compute nodes. The Nexus Dashboard Fabric Controller application framework takes care of the complete life cycle management of all applications that run on the compute nodes. The EPL instance runs as a container that has its own IP address allocated out of the inband pool assigned to the compute nodes. This IP address will be in the same IP subnet as the one allocated to the eth2 or inband interface. Using this IP address, the EPL instance forms a BGP peering with the spines/RRs when the EPL feature is enabled. If a compute node hosting the EPL instance will go down, the EPL instance will be automatically respawned on one of the remaining 2 compute nodes. All IP addresses and other properties associated with the EPL instance are retained.

The Layer-2 adjacency requirement of the compute nodes dictates that the compute node eth2 interfaces should be part of the same IP subnet as the Nexus Dashboard Fabric Controller nodes. Again, in this case, connecting the compute nodes to the same vPC pair of switches is the recommended deployment option. Note that for

cluster mode Nexus Dashboard Fabric Controller OVA setups, ensure that promiscuous mode is enabled in the port group corresponding to eth2 interface in order to establish inband connectivity as depicted below:

EPL-Inband - Edit Settings


Properties			
Security	Promiscuous mode	<input checked="" type="checkbox"/> Override	Accept
Traffic shaping	MAC address changes	<input checked="" type="checkbox"/> Override	Accept
Teaming and failover	Forged transmits	<input checked="" type="checkbox"/> Override	Accept

CANCEL

OK

The enablement of the EPL feature for Nexus Dashboard Fabric Controller cluster mode is identical to that in the non-cluster mode. The main difference is that on the spine/RRs, only a single BGP neighborship is required that points to the IP address allocated to the EPL instance. Recall that for the Nexus Dashboard Fabric Controller native HA deployment in the non-cluster mode, all spines/RRs always had 2 configured BGP neighbors, one pointing to the Nexus Dashboard Fabric Controller primary eth2 interface and other one pointing to the Nexus Dashboard Fabric Controller secondary eth2 interface. However, only one neighbor would be active at any given time.

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  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 **Control > Fabric Builder**.

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.

EPL Connectivity Options

Sample topologies for the various EPL connectivity options are as given below.

Cisco Nexus Dashboard Fabric Controller supports the following web browsers:

NDFC Cluster Mode: Physical Server to VM Mapping

We recommend a minimum of 3 physical servers, or a maximum of 5 physical servers in which each Nexus Dashboard Fabric Controller and compute is located on an individual physical server.

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**.

Troubleshooting Endpoint Locator

There may be multiple reasons why enabling the Endpoint Locator feature may fail. Typically, if the appropriate devices are selected and the IP addresses to be used are correctly specified, the connectivity of the Nexus Dashboard Fabric Controller to the BGP RR may not be present due to which the feature cannot be enabled. This is a sanity check that is present to ensure that basic IP connectivity is available. The following image shows an example error scenario that was encountered during an attempt to enable the EPL feature.

The log that provides details on what occurred when the EPL feature is enabled or disabled, is present in the file `epl.log` at the location: `/usr/local/cisco/dcm/fm/logs/epl.log`. The following example provides a snapshot of the `epl.log` that shows the EPL configuration progress for a fabric.

```
2019.12.05 12:18:23 INFO [epl] Found NDFC Active Inband IP: 192.168.94.55/24
2019.12.05 12:18:23 INFO [epl] Running script: [sudo, /sbin/appmgr, setup, inband-route,
--host, 11.2.0.4]
2019.12.05 12:18:23 INFO [epl] Getting EPL configure progress for fabric 4
2019.12.05 12:18:23 INFO [epl] EPL Progress 2
2019.12.05 12:18:23 INFO [epl] [sudo, /sbin/appmgr, setup, inband-route, --host, 11.2.0.4]
command executed, any errors? No
2019.12.05 12:18:23 INFO [epl] Received response:
2019.12.05 12:18:23 INFO [epl] Validating host route input
2019.12.05 12:18:23 INFO [epl] Done configuring host route
2019.12.05 12:18:23 INFO [epl] Done.
2019.12.05 12:18:23 INFO [epl] Running script: [sudo, /sbin/appmgr, setup, inband-route,
--host, 11.2.0.5]
2019.12.05 12:18:23 INFO [epl] [sudo, /sbin/appmgr, setup, inband-route, --host, 11.2.0.5]
command executed, any errors? No
2019.12.05 12:18:23 INFO [epl] Received response:
2019.12.05 12:18:23 INFO [epl] Validating host route input
2019.12.05 12:18:23 INFO [epl] Done configuring host route
2019.12.05 12:18:23 INFO [epl] Done.
2019.12.05 12:18:23 INFO [epl] Running command: sudo /sbin/appmgr show inband
2019.12.05 12:18:24 INFO [epl] Received response: Physical IP=192.168.94.55/24
Inband GW=192.168.94.1
No IPv6 Inband GW found

2019.12.05 12:18:26 INFO [epl] Call:
http://localhost:35000/afw/apps?imtag=cisco:epl:2.0&fabricid=epl-ex-site, Received
response:
2019.12.05 12:18:26 INFO [epl] Epl started on AFW
```

After the EPL is enabled successfully, all the debug, error, and info logs associated with endpoint information are stored in `/var/afw/applogs/` under the directory for the associated fabric. For example, if EPL is enabled for the **test** fabric, the logs will be in `/var/afw/applogs/epl_cisco_test_afw_log/epl/` starting with filename `afw_bgp.log`. Depending on the scale of the network and the number of endpoint events, the file size will increase. Therefore, there is a restriction on the maximum number and size of `afw_bgp.log`. Up to 10 such files will be stored with each file size of maximum of 100 MB.



Note EPL creates a symlink in this directory inside the docker container, hence it appears broken when accessed natively.

The EPL relies on BGP updates to get endpoint information. In order for this to work, the switch loopback or VTEP interface IP addresses must be discovered on the Nexus Dashboard Fabric Controller for all switches that have endpoints. To validate, navigate to the Cisco Nexus Dashboard Fabric Controller **Web UI > Dashboard > Switch > Interfaces** tab, and verify if the IP address and the prefix associated with the corresponding Layer-3 interfaces (typically loopbacks) are displayed correctly.

In a Cisco Nexus Dashboard Fabric Controller Cluster deployment, if EPL cannot establish BGP peering and the active Nexus Dashboard Fabric Controller is able to ping the loopback IP address of the spine, while the EPL container cannot, it implies that the eth2 port group for Cisco Nexus Dashboard Fabric Controller and its computes does not have Promiscuous mode set to **Accept**. After changing this setting, the container can ping the spine and EPL will establish BGP.

In a large-scale setup, it may take more than 30 seconds (default timer set in Cisco Nexus Dashboard Fabric Controller) to get this information from the switch. If this occurs, the `ssh.read-wait-timeout` property (in the **Settings > Server Settings**) must be changed from 30000 (default) to 60000 or a higher value.

The endpoint data displayed on the dashboard may be slightly inaccurate in a large-scale setup. An approximately 1% accuracy tradeoff is made at higher endpoint counts for performance. If the dashboard greatly differs from what is expected, the validity can be checked with a verifier script that is packaged in Nexus Dashboard Fabric Controller. As root, run the `epl-rt-2.py` script in `/root/packaged-files/scripts/`. This script needs the RR/spine IP and the associated username and password. Note that the `/root/packaged-files/scripts/` directory is read only, so the script needs to be run outside that directory. For example, to run the script for a spine with IP 10.2.0.5, username admin, and password cisco123, run **`/root/packaged-files/scripts/epl-rt-2.py -s 10.2.0.5 -u admin -p cisco123`** while the working directory is `/root/`. If the EPL dashboard still does not display expected numbers and the `epl-rt-2.py` script output differs significantly from the dashboard, please contact tech support.

In cluster mode, BGP is not established between the spines/RRs and Nexus Dashboard Fabric Controller. Check that the **Promiscuous mode** setting for the port group corresponding to the eth2 Nexus Dashboard Fabric Controller interface is set to **Accept**. If a connection is still not established, perform the following steps to check the connectivity between Nexus Dashboard Fabric Controller's BGP client and the spine/RR:

1. Open a shell on the active Nexus Dashboard Fabric Controller and run the following commands:
 - a. `docker service ls`
*Note the ID for the EPL service
 - b. `docker service ps $ID`
*Note the NODE field
 - c. `afw compute list -b`
*Note the HostIp matching the HostName (NODE) from before. This is the compute that the EPL service is currently running on.
2. Open a shell on the compute noted from Step 1 - c and run the following commands:
 - a. `docker container ls`

Note the CONTAINER ID for EPL. If there are multiple EPL containers check the container name to see which one corresponds to which fabric. The naming scheme is `epl_cisco_${FabricName}_afw.`

b. `docker container inspect $CONTAINER_ID`

*Note the value of `SandboxKey`

c. `nsenter --net=$SandboxKey`

This command enters the network namespace of the EPL container. Now network commands such as `ifconfig`, `ip`, and `ping` will act as if they're being ran from inside the container until "exit" is issued in the shell.

3. Try pinging the spine/RR. Make sure that the Inband IP Pool that the Nexus Dashboard Fabric Controller cluster is configured with does not conflict with any switch loopback IPs.

EPL with ISE Policies

Consider a scenario in which AAA configurations are configured on switches running Cisco NX-OS Release 9.3(4) or earlier releases. A sample AAA switch configuration is given below.

```
feature tacacs+
tacacs-server host ISE_ACS_IP_ADDRESS 5 key 7 "Fewhg12345"
aaa group server tacacs+ AAA_TACACS
    server ISE_ACS_IP_ADDRESS
    use-vrf management
    source-interface mgmt0
aaa authentication login default group AAA_TACACS local
aaa authentication login console local
aaa authorization config-commands default group AAA_TACACS local
aaa authorization commands default group AAA_TACACS local
aaa accounting default group AAA_TACACS
aaa authentication login error-enable
```

The ISE server is configured such that the **guestshell**, **run guestshell**, and **show** commands, are permitted to reach the discovery account or policies that are created in the ISE. The permitted commands are set in the **TACACS Command Sets** window under the **Policy Elements** tab in the ISE.

The screenshot shows the Cisco Identity Services Engine (ISE) web interface. The breadcrumb navigation is: Home > Context Visibility > Operations > Policy > Administration > Work Centers. The main navigation bar includes: Network Access, Guest Access, TrustSec, BYOD, Profiler, Posture, Device Administration, and PassiveID. The sub-navigation bar includes: Overview, Identities, User Identity Groups, Ext Id Sources, Network Resources, Policy Elements (selected), Device Admin Policy Sets, Reports, and Settings.

The main content area is titled "TACACS Command Sets". It features a toolbar with icons for Refresh, Add, Duplicate, Trash, Edit, Import, and Export. Below the toolbar is a table with the following data:


<input type="checkbox"/>	Name	Description
<input type="checkbox"/>	DenyAllCommands	Default Command Set
<input type="checkbox"/>	PermitAll	
<input type="checkbox"/>	dcnm-admins-all-priv	
<input type="checkbox"/>	dcnm-discovery-priv	
<input type="checkbox"/>	nexus-admins-all-Priv	

The eth0 IP of Nexus Dashboard Fabric Controller and the subnet for the fabric devices are also allowed. This is configured in the **Device Admin Policy Sets** window under the **Device Administration** tab.

The screenshot shows the Cisco Nexus Dashboard Fabric Controller interface. The navigation menu on the left includes: Overview, Identities, Id Groups, Ext Id Sources, Network Resources, Policy Elements, Policy Sets, Troubleshoot, Reports, Settings, and Dictionaries. The main content area is titled 'Device Admin Policy Sets' and shows a table of policy sets. Two policy sets are visible:

Name	Conditions	Use	Hits	Actions
Discovery Account Permit	AND <ul style="list-style-type: none"> TACACS User EQUALS domes TACACS Remote-Address STARTS_WITH 10.195.198 TACACS Remote-Address STARTS_WITH 172.29.140 TACACS Remote-Address STARTS_WITH 172.28.168 TACACS Remote-Address STARTS_WITH 192.168.10 TACACS Remote-Address STARTS_WITH 172.28.2 	Internal Users	0	Options
Discovery Account Deny	AND <ul style="list-style-type: none"> TACACS User EQUALS domes TACACS Remote-Address NOT_STARTS_WITH 172.29.140 TACACS Remote-Address NOT_STARTS_WITH 10.195.198 TACACS Remote-Address NOT_STARTS_WITH 172.28.168 TACACS Remote-Address NOT_STARTS_WITH 192.168.10 TACACS Remote-Address NOT_STARTS_WITH 172.28.2 	DenyAccess	0	Options

Now, Nexus Dashboard Fabric Controller is configured to use the discovery account to run all the **show** commands that are required for the Endpoint Locator feature. However, due to an issue with the switch NXAPI, AAA validation fails as the requestor IP is not populated in the remote AAA authorization requests. Since the **show** commands are not seen as originating from an IP address, the commands are blocked which prevents the EPL dashboard from displaying the required endpoint information.

As a workaround, we recommend relaxing AAA rules and allowing requests from "blank" senders. To allow requests from "blank" senders, click the  icon under the **Status** column for both **Discovery Account Permit** and **Discovery Account Deny** in the **Device Admin Policy Sets** window, choose **Disabled**, and click **Save**.

Also, this issue is not seen on switches running Cisco NXOS Release 9.3(5) and later releases.



CHAPTER 6

Switches

- [Switches, on page 221](#)
- [Switch Overview, on page 241](#)

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. Status will be either In-Sync or Out-of-sync.
Oper Status	Specifies the configuration status. 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 via the fabric interface. An operator must add static routes on Nexus Dashboard Management Network to connect with switches that needs to be reached over Management interface (bond1br). This ensures that a route for the pods use Management interface as the exit interface.

To add switches to 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 a 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.

Additionally, you can pre-provision switches and interfaces. For more information, see pre-provision device and pre-provisioning ethernet interface.



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 prior to 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**

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.

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

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

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 the Out-of-Sync, then you must deploy the changes. The process is the same as explained in the *Discovering Existing Switches* section.

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

11. After the pending configurations are deployed, the **Progress** column displays 100% for all switches.
12. Click **Close** to return to the fabric builder topology.
13. Click **Refresh Topology** to view the update. All switches must be in green color indicating that they are functional.
14. The switch and the link are discovered in 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 Easy_Fabric_eBGP 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 [Assigning Set 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.

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 PTL.

Note the following guidelines:

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

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.

- [POAP RMA Flow](#)

- [Manual RMA Flow](#)
- [RMA for User with Local Authentication](#)

Prerequisites

- Ensure that the fabric is up and running with minimal disruption while replacing the switch.
- To use the POAP RMA flow, configure the fabric for bootstrap (POAP).
- Perform save and deploy more than once, if needed, to copy the FEX configurations for the RMA of switches that have FEX deployed.

Guidelines and Limitations

- To replace the switch, remove the old switch from the fabric and discover the new switch in the fabric. For example, if you want to replace a Cisco Nexus 9300-EX switch with a Cisco Nexus 9300-FX switch, remove the 9300-EX switch from the fabric followed by discovering the 9300-FX switch in the same fabric.
- When GIR is enabled before upgrading Cisco Nexus 7000 Series switches, 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](#).

Manual RMA Flow

Use this flow when Bootstrap is not possible (or not desired).

To provision manual RMA, follow below procedure:

Procedure

-
- | | |
|---------------|--|
| Step 1 | Place the device in maintenance mode (optional). |
| Step 2 | Physically replace the device in the network. |
| Step 3 | Log in through Console and set the Management IP and credentials. |
| Step 4 | The Cisco Nexus Dashboard Fabric Controller rediscovers the new device (or you can manually choose Discovery > Rediscover) |
| Step 5 | Deploy the expected configuration using Deploy . |
| Step 6 | Depending on the configuration, if breakout ports or FEX ports are in use, you have to deploy again to completely restore the configuration. |
| Step 7 | After a successful deployment, and the device is “In-Sync,” you must move the device back to Normal Mode. |
-

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**, and choose **Operation > Maintenance Mode**.
 - 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** Initiate the RMA flow. Choose the device, click **Actions**, and choose **Operation > Provision RMA**.
 - Step 5** Set the admin password.
(Optional) You can set a AAA user and password for discovery.
 - Step 6** Select the replacement device.
 - Step 7** Click **Provision RMA**.
-

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.
-

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 its current running configuration, and a config compliance check is performed for the switch to get the **In-Sync** or **Out-of-Sync** status. If the switch is out of sync, the user is provided with a preview of all the configurations running in that particular switch that vary from the intent defined by the user for that respective switch.

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 Set 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.

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**.

For more information on golden backup, refer to [Golden Backup](#).



Note Most of this information is at the fabric level, and may or may not directly impact the proceedings of the switch-level restore.

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 switch(s)**. A confirmation window appears, click **Confirm**

Switch Overview

You can perform below operations, from **Actions** icon on Switch Overview window:

- [Previewing Switches](#)
- [Deploy Configuration](#)
- [Discovery](#)
- [Assigning Set Roles](#)
- [Creating a vPC Setup](#)
- [Performing Actions on Switches](#)

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.

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:

Viewing Information for Switch 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.

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. 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.

Field	Description
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: <ul style="list-style-type: none"> • Creating Inter-Fabric Links, on page 129 • Creating Intra-Fabric Links, on page 127
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. <p>Note</p> <ul style="list-style-type: none"> • You cannot update existing links. • The Import Links icon is disabled for external fabric.
Export	Choose the link and select Export to export the links in a CSV file. <p>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.</p>

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.

Interfaces

This section contains the following topics:

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 Adding a Policy

Action Item	Description
Edit Policy	<p>Choose a policy from the table and choose Edit Policy to modify the policy.</p> <p>Note</p> <ul style="list-style-type: none"> • The policies in the italics font cannot be edited. The value under the Editable and Mark Deleted columns for these policies is false. • A warning appears when you edit a policy whose Mark Deleted value is set to <i>true</i>. 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	<p>Choose policies from the table and choose Delete Policy to delete the policies.</p> <p>Note</p> <p>A warning appears when you delete policies whose Mark Deleted values are set to <i>true</i>.</p>
Generated Config	<p>Choose policies from the table and choose Generated Config to view the delta of configuration changes made by every user.</p>
Push Config	<p>Choose policies from the table and choose Push Config to push the policy configuration to the device.</p> <p>Note</p> <ul style="list-style-type: none"> • This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric. • A warning appears if you push configuration for a Python policy. • A warning appears when you push configurations for policies whose Mark Deleted values are set to <i>true</i>.

Event Analytics

Event Analytics includes the following topics:

- [Alarms, on page 295](#)

- [Events, on page 300](#)
- [Accounting, on page 304](#)

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.
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.

Services

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 service devices. You

can add a service node, create route peering between the service node and the service leaf switch, and then selectively redirect traffic to these service nodes.



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.



CHAPTER 7

Policies

- [Viewing and Editing Policies, on page 251](#)
- [Adding a Policy, on page 253](#)

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 Adding a Policy
Edit Policy	<p>Choose a policy from the table and choose Edit Policy to modify the policy.</p> <p>Note</p> <ul style="list-style-type: none"> The policies in the italics font cannot be edited. The value under the Editable and Mark Deleted columns for these policies is false. A warning appears when you edit a policy whose Mark Deleted value is set to <i>true</i>. 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	<p>Choose policies from the table and choose Delete Policy to delete the policies.</p> <p>Note A warning appears when you delete policies whose Mark Deleted values are set to <i>true</i>.</p>
Generated Config	Choose policies from the table and choose Generated Config to view the delta of configuration changes made by every user.

Action Item	Description
Push Config	<p>Choose policies from the table and choose Push Config to push the policy configuration to the device.</p> <p>Note</p> <ul style="list-style-type: none">• This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.• A warning appears if you push configuration for a Python policy.• A warning appears when you push configurations for policies whose Mark Deleted values are set to <i>true</i>.

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**.
- Step 3** Click **Choose Template** and choose appropriate policy template and click **Select**.
- Step 4** Enter the mandatory parameters in the text field and click **Save**.
-



CHAPTER 8

Interfaces

This section contains the following topics:

- [Interfaces, on page 255](#)
- [Interface Groups, on page 265](#)

Interfaces

The Interfaces option displays all the interfaces that are discovered for the switch, Virtual Port Channels (vPCs), and intended interfaces missing on the device.

You can use the following functions:

- Create, deploy, view, edit and delete a port channel, vPC, Straight-through FEX, Active-Active FEX, loopback, and subinterface.



Note

- The following features are unsupported for the brownfield migration of switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images:
 - FEX on switches other than Cisco Nexus 9300 Series switches and Cisco Nexus 9500 Series switches with X9500 line cards
 - AA-FEX

For information about the platform support for FEX, refer to your platform and NX-OS documentation to check the feature compatibility.

- To edit interfaces associated with fabric links such as intra-fabric links and inter-fabric links, see [Editing Interfaces Associated with Links, on page 262](#).
 - 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 new interface	Allows you to add a logical interface such as a port channel, vPC, Straight-through FEX, Active-Active FEX, and loopback. For more information, see Adding Interfaces, on page 259 .
Create new subinterface	Allows you to add a logical subinterface.
Edit interface	Allows you to edit and change policies that are associated with an interface.
Preview interfaces	Allows you to preview the interface configuration.

Field	Description
Deploy interfaces	Allows you to deploy or redeploy saved interface configurations.
Shutdown	Allows you to shut down the interface.
No Shutdown	Allows you to enable an interface (no shutdown or admin up).
Breakout	Allows you to <i>breakout</i> an interface.
Un-Breakout	Allows you to unbreakout interfaces that are in <i>breakout</i> state.
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.
Show commands	Allows you to display the interface show commands. A show command requires show templates in the template library.
Rediscover Interface	Allows you to rediscover or recalculate the compliance status on the selected interfaces.
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.
Deployer History	Allows you to display the interface deployment history details.

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

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
Rediscover	Supported	Supported
Deploy	Blocked	Blocked

The buttons for the associated operations are grayed out accordingly.

If you perform admin operations (shutdown/no shutdown) on SVI, which is part of a config profile, successive **Save & Deploy** operations generate **no interface vlan** command.

For SVI with no policy, on performing admin operation, that is, shutdown/no shutdown command pushed from **Interface Manager, int_vlan_admin_state** policy is associated with the SVI.

For example, create and deploy the SVI from **switch_freeform**.

```
interface vlan1234
  description test
  no shutdown
  no ip redirects
  no ipv6 redirects
```

If you shutdown the SVI from interface manager, the **int_vlan_admin_state** policy is associated with the SVI.

Pending diff is shown as:

```
interface Vlan1234
  shutdown
  no ip redirects
  no ipv6 redirects
  description test
  no shutdown
```

Remove the **no shutdown** CLI from the free-form config.

If the user has performed admin operation on SVI, device will have interface in running config. Therefore, post network detach **interface vlan** will be still present and interface will be discovered. You need to manually delete the interface from **Interface Manager**.

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.
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.

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. External fabric devices aren't listed for ST FEX and AA FEX. In the case of vPC or Active to Active FEX, select the vPC switch pair.
- Step 5** Enter the ID value in the respective interface ID field (**Port Channel ID**, **vPC ID**, **Loopback ID**, **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.

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

Click the drop-down arrow next to the **Breakout** icon to display a list of the available breakout options. The available options are **10g-4x**, **25g-4x**, **50g-2x**, **50g-4x**, **100g-2x**, **100g-4x**, **200g-2x**, and **Unbreakout**. Choose the required option.

Editing Interfaces

To edit the interfaces from the Cisco 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 or deploy the changes.
 - Step 4** Click **No Shutdown** to bring up the selected interfaces.
A confirmation window appears where you can save, preview, and deploy the changes. Click **Save** to preview or deploy the changes.
-

Viewing Interface Configuration

To view the interface configuration commands and execute them from the Cisco 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

You can create an interface group that allows grouping of host-facing interfaces at a fabric level. Specifically, you can create an interface group for physical Ethernet interfaces, L2 port-channels, and vPCs. You can attach or unattach multiple overlay networks to the interfaces in an interface group.

Guidelines

- Interface groups are only supported for the fabrics with the **Easy_Fabric** template.
- An interface group is specific to a fabric. For example, consider two fabrics: Fab1 and Fab 2. The interface group IG1 in Fab1 isn't applicable to Fab 2.
- An interface group can only have interfaces of a certain type. For example, you need three separate interface groups if you want to group three types of interfaces such as IG1 for physical Ethernet trunk interfaces, IG2 for L2 trunk port-channels, and IG3 for vPC host trunk ports.
- An interface group can also be created using preprovisioned interfaces.
- Interface groups are limited to switches with the leaf role. They aren't supported for other roles such as Border, BGW, and other related variants.
- For L2 port-channels and vPCs that are part of an interface group, they can't be deleted until they are de-associated from the interface group even if there are no networks associated with the interface group. Similarly, a trunk port that has no overlay networks but is part of an IG can't be converted to an access port. In other words, you can't change policies for interfaces that are part of an interface group. However, you can edit certain fields for policies.
- For L4-L7 services configuration on leaf switches, trunk ports that are used for services attachment can't be part of interface groups.
- When you perform a per fabric backup of an easy fabric, if there are interface groups created in that fabric, all the associated interface group state is backed up.
- If an easy fabric contains an interface group, then this fabric can't be imported into the MSO. Similarly, if an easy fabric has been added to the MSO, you can't create interface groups for interfaces that belong to switches in the easy fabric.
- The **Interface Group** button is enabled only for Admin and Stager users. For all other users, this button is disabled.
- The **Interface Group** button is disabled in the following circumstances:
 - Select **Data center** from the **SCOPE** drop-down list.
 - Select a fabric without any switches.
 - Select any other interface apart from vPC, Port-channel, and Ethernet.
 - If the interface has a policy attached from another source, for example:
 - If the interface is member of a port-channel or vPC.
 - If the port-channel is member of vPC.
 - If the interface has a policy from underlay or links.



Note If you select different types of interfaces, the **Interface Group** button is enabled. However, when you try to create or save different types of interfaces to an interface group, an error is displayed.

Creating an Interface Group

Procedure

- Step 1** 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 belong to only a single interface group.
- Step 5** Click the Interfaces tab.
- Step 6** Select the interfaces that have to be grouped and click **Actions > Add to interface Group**.
- Step 7** In the **Edit Interface Group** window, create a custom interface group by entering an interface group name in the **Select Interface Group** field and click **Create custom**.
- 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 8** Click **Save**.
- In the **Interfaces** window, you can see the interface group name under the **Interface Group** column.
-

Removing Interfaces from an Interface Group

Procedure

- Step 1** Choose **LAN > Interfaces**.
- Step 2** Select the interfaces to disassociate from an interface group and click **Actions > Remove from interface Group**.
- Step 3** In the **Edit Interface Group** window, make sure that nothing is selected in the **Select Interface Group** drop-down list, and click **Clear**.
- A dialog box pops up asking whether you want to clear all the associated interfaces. Click **Yes** to proceed. Note that if there are any networks attached to these interfaces, they are detached as well when you click **Clear**.
-

Attaching Networks to an Interface Group

Procedure

Step 1 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 xxx**” 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 custom**. 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 267](#).

Step 4 Click **Continue** and click **Save & Deploy** to deploy the selected networks on the switches.

Unattaching a Network from an Interface Group

This procedure shows how to unattach a network from an interface group in the Networks window. Also, you can unattach networks when you remove an interface from an interface group in the **Interfaces** window. For more information, see *Removing Interfaces from an Interface Group*.

Procedure

- Step 1** 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**.
 - Step 3** In the **Interface Groups** window, select the interface group from the **Select Interface Group** drop-down list and click **Clear** to unattach a network.
 - Step 4** (Optional) Navigate to **LAN > Interfaces**.

Under the **Overlay Network** column, you can see the unattached network in the red color for the corresponding interface. Click the network to view the expected config that is struck through.
 - Step 5** Navigate to the **Networks** screen. From Fabrics Actions drop-down list, select **Deploy Config**.
-

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. 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**.
-



CHAPTER 9

Services

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 service devices. You can add a service node, create route peering between the service node and the service leaf switch, and then selectively redirect traffic to these service nodes.



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.

- [Services, on page 271](#)

Services

Cisco Nexus Dashboard Fabric Controller provides ability to insert service devices in a data center fabric, and also enables selectively redirecting traffic to these service devices. You can add a service node, create route peering between the service node and the service leaf switch, and then selectively redirect traffic to these service nodes.

You can also watch a video that demonstrates how to orchestrate a 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 Node

You have to create an external fabric and specify that a service node resides in that external fabric 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, or a border super spine. Nexus Dashboard Fabric Controller does not define a new switch role for a service leaf.

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 service will not change its mode. In case the attached switches are forming a vPC pair, the name of the attached switch is a combination of both switches.

Route Peering

Route peering creates service networks. 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.

The automatically created service network will also be listed on the **LAN > Services > Service Nodes** window. However, you cannot delete the service network. Deletion of service networks is handled automatically during the service route peering deletion process. There can be multiple route peerings defined per tenant/VRF.

Service Policy

You can 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 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 **LAN > Services > Service Nodes** window. For intra-tenant firewall, one-arm and two-arm load balancer, the 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 service integration is built on top of the easy fabric policy enforcement. Use the fabric builder to create a VXLAN EVPN fabric and then import Cisco Nexus 9000 Series switches into the fabric with predefined fabric policies.

MSD Support

This feature supports Multi-Site Domains (MSD). Select the MSD member fabric from the Nexus Dashboard Fabric Controller fabric scope selector, create a service node (for example, firewall, or load balancer), attach the service node to the switch in the selected MSD member fabric, define the route peering and service policies, and deploy relevant configurations on the selected MSD member fabric. .

RBAC Support

The 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



Note If a fabric is in fabric monitor or read-only mode, an admin cannot deploy the route peering or service policy. Also, the icon to delete the service node is not displayed if the external fabric where the service node is located is in fabric Monitor Mode. Remove the fabric from the fabric Monitor Mode to display the icon to delete the service node. This icon will be shown only to users with admin role access.

The Service windows are displayed based on the logged-in user's role and reflect the actions that the user is allowed to perform. Example screenshots of the Service Nodes window for an admin, stager, and operator role are as given below:

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. DCNM filters out the subinterfaces of Layer 3 port-channel as PBR is not supported with Layer 3 port-channel subinterfaces. 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 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 Services

- Service in Nexus Dashboard Fabric Controller does not manage or provision service nodes, such as firewall and load balancer.
- The Service feature is supported only on the VXLAN BGP EVPN fabrics with the **Easy_Fabric_11_1** 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 and two-arm deployed load balancers, 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.
- One-arm Virtual Network Function is supported.
- 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.
- 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 > Services** window.

Types of Service Devices

The 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 service attachment is generic and applies to any vendors service node.

Configuring Fabric Settings for Service

Certain fabric settings have to be configured to enable 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.

Now, click **Resources**. Specify a VLAN range in the **Service Network VLAN Range** field. This is a per switch overlay service network VLAN range. The minimum allowed value is 2 and the maximum allowed value is 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 Services

To launch the Services, or the Elastic Service, on the Cisco Nexus Dashboard Fabric Controller Web UI, choose **LAN > Services**.

The **Service Nodes** window is displayed. Select a valid switch fabric to display or define the service nodes, route peerings, and service policies, in that fabric.



Note Service nodes, route peering, and service policies updated within the last 15 minutes are highlighted.

The services configuration procedure consists of the following steps:

Adding Service Node

To add a service node, click **Actions > Add** at the top right of the **Service Nodes** window to display the **Create New Service Node** window.

The **Create New Service Node** window has three 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.

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 fabric 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 Leaf 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. .

A form is displayed depending on the template used. Update all the required fields in the form and click **Save**.

Creating Route Peering

The fields that appear in the **Create Route Peering** window depend on the type of deployment chosen in the **Create New Service Node** window. Depending on the type chosen (Firewall or Load Balancer), the types of deployments are Intra-Tenant Firewall, Inter-Tenant Firewall, One-Arm load balancer and Two-Arm load balancer.



Note Deletion of service network is not supported in Top-down provisioning.

Deletion of service network is not allowed on the **LAN > Services > Service Nodes** 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 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 280](#).

Outside Network

VRF: Specify the VRF.

Network Type: Select Outside Network.

Service Network: Specify the name of the service network.

Vlan ID: Specify the VLAN ID. Valid IDs range from 2 to 3967. Click Propose to retrieve a value from the pre-defined service network VLAN range pool.

Service Network Template: Select the `Service_Network_Universal` template from the drop-down list. For more information on the template fields, refer [Templates, on page 280](#).

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 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 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 280](#).

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 280](#).

Outside Network

VRF: Select a VRF from the drop-down list..

Network Type: Select Outside Network.

Service Network: Provide a 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 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 280](#).

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 280](#).

Example: One-Arm Mode Load Balancer

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

Peering Name: Specify a name for the peering. The name can have alphanumeric, underscore, or dash characters.

Deployment: Select One-Arm Mode.

Peering Option: Select Static Peering or eBGP Dynamic Peering.

First Arm

VRF: Select a VRF from the drop-down list..

Network Type: Select First Arm.

Service Network: Provide a 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 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 280](#).

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 280](#).

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 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 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 280](#).

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 280](#).

Second Arm

VRF: Select a VRF from the drop-down list..

Network Type: Select Second Arm.

Service Network: Provide a 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 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 280](#).

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 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 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

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 a peering option from the drop-down list.

Source VRF Name: Select a source VRF from the drop-down list.

Destination VRF Name: Select a destination VRF from the drop-down list.

Source Network: Select an IP address from the drop-down list.

Destination Network: Select an IP address from the drop-down list.

Reverse Next Hop IP : The reverse next-hop IP address is displayed.

Link Template : Select a template from the drop-down list. For more information on the template fields, refer [Templates, on page 280](#).

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 has been introduced. The options in this 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 280](#).

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.

Deleting a Service Node

To delete a service node from the Cisco DCNM 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.


Editing a Service Node

To edit a service node from the Cisco DCNM Web UI, perform the following steps:

Procedure

- Step 1** Select a service node from the table and click **Edit** or **Actions > Edit**.
- Step 2** The **Edit Service Node** window is displayed.
- Make the required changes and click **Save**.
-

Refreshing a Service Node

To refresh the list of service node that is displayed in the **Service Nodes** window, click the **Refresh** icon .

Service Node Backup and Restore

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 restore the service node level data by clicking **Actions > Import** to import data about the service nodes from an excel file.

You can also export data for a specific service node by selecting the node and clicking **Actions > Export**.

Importing Service Nodes

To import service nodes from as an Excel file, click **Actions > Import** on the **Service Nodes** window. Click the **Import** icon on the **Service Node Import** window to import information about the service nodes.

Exporting Service Nodes

To export a service node as an Excel file, select a service node from the table and click **Actions > Export**. Click **Export** on the **Service Node Export** windows to export information about the service nodes.

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.

The Audit Logs table in the Audit History window displays information about all the actions that have been performed. Audit logs are generated when the following actions are performed:

- 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

This 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.

Information such as the name of the User Name, User Role, Action Taken, Entity, Details, Status, and Time of execution is displayed. Click **Actions > Purge Audit History** to delete the audit history.



PART **III**

Settings

- [Server Settings, on page 287](#)
- [Feature Management, on page 289](#)
- [Credentials Management, on page 291](#)



CHAPTER 10

Server Settings

- [Server Settings, on page 287](#)

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.



CHAPTER 11

Feature Management

- [Feature Management, on page 289](#)

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 one of the Feature set to view the default set of features in the table below.
- Step 3** In the table below, select the applications available with the feature set.
- Step 4** Click **Save & Continue**.
- 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.
-

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 39: 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
SAN Controller	Not supported	Not supported	-



CHAPTER 12

Credentials Management

- [LAN Credentials Management, on page 291](#)

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.
- **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 Management allows you to specify configuration change credentials. Before changing any LAN switch configuration, you must furnish *Configuration Change* SSH 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	Displays the encrypted form of the SSH password.
Username	Specifies the username of the switch Nexus Dashboard Fabric Controller user.
Fabric	Displays the fabric to which the switch belongs.



PART **IV**

Operations

- [Event Analytics, on page 295](#)
- [Image Management, on page 305](#)
- [Programmable Reports, on page 317](#)
- [License Management, on page 323](#)
- [Templates, on page 333](#)
- [Tech Support, on page 363](#)
- [Backup and Restore, on page 365](#)
- [NXAPI Certificates, on page 369](#)



CHAPTER 13

Event Analytics

This section contains the following topics:

- [Alarms, on page 295](#)
- [Events, on page 300](#)
- [Accounting, on page 304](#)

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

After you create a new alarm policy, navigate to **Alarms Raised** tab, click **Refresh** icon to view the created alarm.

Click on 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.
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 **Operations > Event Analytics > Alarms > Alarms Raised**.

Action Item	Description
Acknowledge	Select alarm and choose Acknowledge to apply the alarm
Unacknowledge	Select alarm and choose Unacknowledge to deny the alarm policy
Clear	Select alarm and choose Clear to apply the alarm policy
Delete Alarm	Select an alarm and choose Delete to delete the alarm

Alarms Cleared

This tab displays the cleared alarms. 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 select one or more alarms and then click the Delete button to delete them.

The following table describes the fields that appear on **Operations > Event Analytics > Alarms > Alarms Cleared**.

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.
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 **Operations > Event Analytics > Alarms > Alarms Cleared**.

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: <ul style="list-style-type: none"> • Activated • Deactivated
Policy type	Specifies the type of the policy: <ul style="list-style-type: none"> • Device Health Policy • Interface Health Policy • Syslog Alarm 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.

Action Item	Description
Import	Select to import alarm policies in bulk from a .csv file.
Export	Select to export alarm policies in bulk from a .csv 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.

Create new alarm policy

You can add alarm policies for the following:

- Device Health Policy
- Interface Health Policy
- Syslog Alarm Policy

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.

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. Under **Device Features**, you can select the BFD, BGP, and HSRP protocols. When these check boxes are selected, alarms are triggered for the following traps: **BFD**- ciscoBfdSessDown, ciscoBfdSessUp, **BGP**- bgpEstablishedNotification, bgpBackwardTransNotification, cbgpPeer2BackwardTransition (), cbgpPeer2EstablishedNotification, and **HSRP**- cHsrpStateChange. Please refer <https://snmp.cloudapps.cisco.com/Support/SNMP/do/BrowseOID.do?local=en> for detailed trap OID definition.

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.

Select the devices for which you want to create policies. Specify the policy name, description, link-state, Bandwidth (In/Out), Inbound errors, Outbound errors, Inbound Discards, and Outbound Discards.

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 40: 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 41: 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 42: Example 3:

Identifier	ID1-ID2
Raise Regex	ETHPORT-5-IF_SFP_WARNING: Interface \$(ID1), High Rx Power Warning
Clear Regex	ETHPORT-5-IF_SFP_WARNING: Interface \$(ID1), High Rx Power Warning cleared

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. 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. 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.
Event Setup	Allows you to setup new event. For more information, see Event Setup, on page 301.

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 > Event Setup**. From the **Actions** menu drop-down list, choose **Event Setup**.
- Step 2** In the Receiver tab, perform the following steps:
- Use the toggle button to enable this feature.
 - Select **Copy Syslog Messages to DB** and click **Apply** to copy the syslog messages to the database. If this option is not selected, the events will not be displayed in the events page of the Web client. The columns in the second table display the following:
 - Switches sending traps
 - Switches sending syslog
 - Switches sending syslog accounting
 - Switches sending delayed traps
 - In the Sources tab, the table displays fabrics and switches associate with it. It also displays information about traps and syslogs.
- Step 3** To add and remove notification forwarding for system messages from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:
- Cisco Nexus Dashboard Fabric Controller Web UI forwards fabric events through email or SNMPv1 traps. Some SMTP servers may require addition of authentication parameters to emails that are sent from Nexus Dashboard Fabric Controller to the SMTP servers. You can add authentication parameters to the emails that

are sent by Nexus Dashboard Fabric Controller to any SMTP server that requires authentication. Enable this feature on **Settings > Server Settings > Events** tab.

- a) Choose **Settings > Server Settings > Events** tab. Check **Enable Event forwarding** check box to enable events forwarding. The events forwarding scope, the recipient email address, severity of the event and type of the event is displayed. The description Regex field is applicable only when the forwarding source is selected as Syslog while adding the events forwarder.
- b) Specify the **SMTP Server** details and the **From** email address. Configure the **Snooze** and **Event Count** filter.
- c) Click **Save**.
- d) Choose **Operations > Event Analytics**. From the Actions drop-down list, choose **Add Rule**.
- e) In the Forwarding Method, choose either **E-mail** or **Trap**.

If you choose **Trap**, **Address** and **Port** field is added to the dialog box.

- f) If you choose the **E-mail** forwarding method, enter the IP address in the Email Address field. If you choose the **Trap** method, enter the trap receiver IP address in the Address field and specify the port number.

You can either enter an IPv4 or IPv6 addresses or DNS server name in the Address field.

- g) In the **Fabric** field, choose all groups or specific fabric for notification. For SAN Installer, select **VSAN Scope**. You can either choose **All** or **List** option. If you select List, provide the list of VSANs for notification.
- h) In the Source field, select Nexus Dashboard Fabric Controller or **Syslog**.
 - If you select Nexus Dashboard Fabric Controller, then:
 1. From the **Type** drop-down list, choose an event type.
 2. Check the **Storage Ports Only** check box to select only the storage ports.
 3. From the **Minimum Severity** drop-down list, select the severity level of the messages to receive.
 4. Click **Add** to add the notification.
 - If you select Syslog, then:
 1. In the **Facility** list, select the syslog facility.
 2. Specify the syslog **Type**.
 3. In the **Description Regex** field, specify a description that matches with the event description.
 4. From the **Minimum Severity** drop-down list, select the severity level of the messages to receive.
 5. Click **Add** to add the notification.

Note The Minimum Severity option is available only if the Event Type is set to **All**.

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
```

i) Click **Add Rule**.

Step 4

To add rules to the Event Suppression from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

Cisco Nexus Dashboard Fabric Controller allows you to suppress the specified events that are based on the user-specified suppressor rules. Such events will not be displayed on the Cisco Nexus Dashboard Fabric Controller Web UI and SAN Client. The events will neither be persisted to Nexus Dashboard Fabric Controller database, nor forwarded via email or SNMP trap.

You can view, add, modify, and delete suppressor rules from the table. You can create a suppressor rule from the existing event table. Select a given event as the template and invoke the rule dialog window. Event details are automatically ported from the selected event in the event table to the input fields of the rule creation dialog window.

Note You cannot suppress EMC Call Home events from the Cisco Nexus Dashboard Fabric Controller Web UI.

- a) Specify the **Name** for the rule.
- b) Select the required **Scope** for the rule that is based on the event source.

In the **Scope** drop-down list, the LAN groups and the port groups are listed separately. You can choose **SAN/LAN, Port Groups** or **Any**. 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 suppressor rule is applied globally.

- c) Enter the **Facility** name or choose from the SAN/LAN Switch Event Facility List.

If you do not specify a facility, wildcard is applied.

- d) From the drop-down list, select 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** box and select a valid time range during which the event is suppressed.

By default, the time range is not enabled, i.e., the rule is always active.

Note In general, you must not suppress accounting events. Suppressor rule for Accounting events can be created only for certain rare situations where Accounting events are generated by actions of Nexus Dashboard Fabric Controller or switch software. For example, lots of '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 the Suppressor table and invoke the Add Event Suppressor Rule dialog window.

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.



CHAPTER 14

Image Management

- [Image Management, on page 305](#)

Image Management

Upgrading your devices to the latest software version manually might take a long time and prone to error, which requires a separate maintenance window. To ensure rapid and reliable software upgrades, image management automates the steps associated with upgrade planning, scheduling, downloading, and monitoring. Image management is supported only for Cisco Nexus switches.



Note Before you upgrade, ensure that the POAP boot mode is disabled for Cisco Nexus 9000 Series switches and Cisco Nexus 3000 Series switches. To disable POAP, run the `no boot poap enable` command on the switch console. You can however, enable it after the upgrade.

The **Image Management** window has the following tabs.

Tabs	Actions
Overview	You can perform below actions on Overview tab
	Stage Image
	Validate
	Upgrade
	Change Mode
	Apply Policy
	Recalculate Compliance
Run Reports	
Images	Upload
	Delete
Image Policies	Create
	Delete

Tabs	Actions
History	

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

The Overview window displays all the switches you discovered 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.

Nexus Dashboard Fabric Controller UI Navigation

- Choose **Operations > Image Management > Overview**. Click Actions to perform various operations.

You can perform the following actions in the Overview window:

- Staging an Image
- Validating an Image
- Upgrading an Image
- Apply Policy

- Recalculating Compliance
- Running a Report

Based on the actions you perform, the value under the Reason column is updated.

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:

Procedure

- Step 1** Choose **Operations > Image Management > Overview**.
- 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'll return to the Overview 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 > Overview**.
- 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, uninstall, or set boot variable for a switch.

Upgrade Options for NX-OS Switches

- **Disruptive**: Choose this option for disruptive upgrades.
- **Allow Non-disruptive**: Choose this option to prevent 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 > Overview**.
- 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:
- Choose an upgrade option from the drop-down list based on how you want to upgrade.
 - (Optional) Check the BIOS Force check box.
You can view the validation status of all the devices.
 - Check the **Golden** check box to perform a golden upgrade.
 - 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.

Change the Mode

You can change the mode of the device. To change the mode of a device from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

Procedure

Step 1 Choose **Operations > Image Management > Overview**.

Step 2 Choose the switch for which you want to change the mode by checking the check box.

Note You can choose more than one switch.

Step 3 Click **Actions > Change Mode**.

The **Change Mode** dialog box appears.

Step 4 Choose a mode from the drop-down list.

Valid options are **Normal** and **Maintenance**.

Step 5 Click **Save and Deploy Now** or **Save and Deploy Later**

You will return to the Overview tab under the Image Management window.

Changing 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 > Overview**.

Step 2 Choose a switch by checking the check box.

Step 3 Click **Actions** and choose **Apply 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 the 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 > Overview**.
- 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.
-

Run Reports

Choose **Reports > Report Definitions**.

Select the checkbox next to the report that has to be generated again. From the **Actions** drop-down list, select **Re-run Report** to run a report job again. A pop-up window is displayed indicating that the report job has been run again.

You can use the **Re-run Report** to generate a report before the scheduled execution time. In case of an **On-demand job**, click **Re-run Report** to generate the report.

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	<p>Specifies the name of the platform. Images, RPMs, or SMUs are categorized as follows:</p> <ul style="list-style-type: none"> • N9K/N3k • N6K • N7K • N77K • N5K • Other • Third Party <p>The images are the same for N9K and N3K platforms.</p> <p>The platform will be Other if the uploaded images are not mapped to any of the existing platforms.</p> <p>The platform will be Third Party for RPMs.</p>
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	<p>Specifies the file type of the image, EPLD, RPM, or SMU.</p> <p>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.</p>
NXOS Version	Specifies the NXOS image version for only Cisco switches.
Image Version	Specifies the image version for all devices, including the non-Cisco devices as well.
Size (Bytes)	Specifies the size of the image, RPM, or SMU files in bytes.
Checksum	<p>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.</p>

The following table describes the action items, in the **Actions** menu drop-down list, that appear 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 312 .
Delete	Allows you to image from the repository. Choose an image, click Actions , and choose Delete . A confirmation window appears. Click Yes to delete the image.

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 **MDS**, **Third Party** or **Other** platform. The **Third Party** platform is applicable only for RPMs.

Step 5 Click **Verify**.

The EPLD images, RPMs, and SMUs are uploaded to the repository in the following path:
`/var/lib/dcnm/upload/<platform_name>`.

All uploaded images copied into MinIo repository under `imageMgmt/`

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>`

All uploaded images copied into MinIo repository under `imageMgmt/`

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.

You can view the policy details in the Image Policies window, and perform the following tasks from this window:

Creating an Image Policy

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



Note While creating a policy for MDS platform and SAN deployment few fields are grayed out.

Before you begin

Upload the images under the **Images** tab before creating an image policy. See the [Uploading an Image, on page 312](#) 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.
Release	Choose the NX-OS version from the Release drop-down list. The release versions of 64-bit images are appended with 64bit in the image name.
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 [Changing a Policy, on page 309](#) section for more information.

Deleting an Image Policy

To delete an image policy from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

Before you begin

Detach policy from the device before deleting it.

Procedure

Step 1 Choose **Operations > Image Management > Image Policies**.

Step 2 Choose a policy by checking the check box.

Step 3 Click **Actions** and choose **Delete**.

A confirmation dialog appears.

Step 4 Click **Confirm**.

Note An error message appears if you try to delete a policy that is attached to a device.

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.

ID	
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 15

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 360](#). 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 to Nexus Dashboard Fabric Controller Release 12.0.1a. You must create them again manually.

This chapter contains the following sections:

- [Report Templates, on page 318](#)
- [Creating a Report Job, on page 318](#)
- [Report Definitions, on page 319](#)
- [Reports, on page 321](#)

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 *Template Library*. Templates are listed based on the associated tags.

Inventory_Report, **Performance_Report**, and **Switch_Performance_Report** are used for performance management reports.

Creating a Report Job

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**.
Base don 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 options are:

Option	Description
Now	The report is generated now.

Option	Description
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.

Note When you are creating a Periodic NVE VNI Counters report, the report generation frequency has to be set to 60 minutes or more. If the frequency is less than 60 minutes, an error message is displayed.

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 is created and appears on the **Reports** tab.

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: <ul style="list-style-type: none"> • Success: Report is generated successfully. • Scheduled: A report generating schedule is set. • Running: A report job is running. • Failed: Report execution failed for one or more selected switches/fabrics or an issue occurred during running of the report job. • Unknown: Job state could not be identified.
Scheduled 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. Note 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	<p>Allows you to view report job history.</p> <p>The Job History window is displayed. You can view several entries per report job.</p> <p>Note 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.</p> <ul style="list-style-type: none"> • Max number of history across report definition • 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	<p>Specifies the title of the report.</p> <ul style="list-style-type: none"> • Single click on the report title opens a slide in summary panel. • 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.

Field	Description
Status	Specifies the status of the report. The status messages are as follows: <ul style="list-style-type: none"> • COMPLETED • SUCCESS • RUNNING • FAILED • WARNING • SCHEDULED • 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. Note 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. 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 16

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 modify existing server license files to smart licenses on Cisco Smart Software Manager (CSSM).

This chapter contains the following topics:

- [Overview, on page 323](#)
- [NDFC Server Licenses, on page 324](#)
- [Smart Licensing, on page 325](#)
- [Switch Licenses, on page 328](#)
- [Switch License Files, on page 329](#)

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 Licensing](#) section.

NDFC Server Licenses

On NDFC tab, you can assign either a switch based smart or a server based smart licenses to the switches.

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: <ul style="list-style-type: none"> • Switch • Smart • Switch Smart
State	Displays the license type of the switch that can be one of the following: <ul style="list-style-type: none"> • Permanent • Unlicensed • Smart • Expired • Not Applicable • Invalid
Expiration Date	Specifies the expiration date of license.
WWN/Chassis ID	Displays the world wide name or Chassis ID.
Model	Displays the model of the device. For example, DS-C9124 or N5K-C5020P-BF.

Field	Description
Fabric	Specifies the name of the fabric.

To add license to

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 . A confirmation message appears.
Unassign	Choose a switch, from the Actions drop-down list, select UnAssign . A confirmation message appears.
Assign All	<ul style="list-style-type: none"> To assign license to all switches in the table, from the Actions drop-down list, choose Assign All. <p>A confirmation message appears</p> <ul style="list-style-type: none"> Click OK to refresh table.
Unassign All	<ul style="list-style-type: none"> To unassign license to all switches in the table, from the Actions drop-down list, choose UnAssign All. <p>A confirmation message appears</p> <ul style="list-style-type: none"> Click OK to refresh table.

Smart Licensing

Cisco Nexus Dashboard Fabric Controller allows you to configure Smart Licensing for

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 (software.cisco.com). For a more detailed overview on Cisco Licensing, go to cisco.com/go/licensingguide.

Smart License Management

The **Smart** page shows the following cards:

- **Enable Smart Licensing**

Use the toggle switch to enable Smart Licensing.

- **Trust Status**

Click on **Establish Trust** to establish trust.

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 **Proxy** to transport using the proxy server. Enter the URL and Port details to access via the proxy server.

Enter the Registration token obtained from CSSM. For instructions, refer to [Obtaining Token from CSSM, on page 327](#).



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.

- **License Status**

Specifies the status of the license. 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 **AUTHORIZED** or **OUT-OF-COMPLIANCE** 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.

CSSM allows you to convert traditional licenses to Smart Licenses. For instructions, refer to [Converting Classic Licenses to Smart Licenses](#).

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 CSSM using SLP. For instructions, refer to [Smart Licensing using Policy to Establish Trust with CSSM, on page 327](#).

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 Authorized and Out of Compliance .
Description	Specifies the type and details of the license.

Obtaining Token from CSSM

To establish trust, you must obtain and enter a valid Registration Token in Cisco Nexus Dashboard Fabric Controller Web UI. To obtain a token from CSSM, perform the following steps:

Procedure

-
- Step 1** Log in to **Cisco Smart Software Manager (CSSM)**.
- Step 2** Choose **Cisco Software Central > Smart Licensing > Inventory** tab.
- Step 3** In **Product Instance Registration Token**, generate a new token.
- A token is required to register product instances, to use licenses from the virtual account.
- Step 4** In **Tokens** table, click on the correct token, and copy it to your clipboard.
- This token is required during Smart License Trust Establishment on **Operations > License Management > Smart** tab on the Cisco Nexus Dashboard Fabric Controller Web UI.
-

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 the **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. For instructions, see [Obtaining Token from CSSM, on page 327](#).

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>.
 - **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 switches will be registered with Smart License.
- Step 9** Click **Smart** tab.
The **Status** of the server licenses shows **InCompliance**.
If the status shows **OutofCompliance**, visit the CSSM portal to acquire the required licenses.
For all other statuses, contact Cisco Technical Assistance Center (TAC).

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 switches 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. <ul style="list-style-type: none"> • Unused • In Use • Out Of Compliance
Type	Displays the license type of the switch that can be one of the following: <ul style="list-style-type: none"> • Temporary • Permanent • Smart • Counter Permanent • Unlicensed • Counted
Warnings	Specifies the warnings about license, such as expiration date and time.
Group	Specifies the fabric or LAN name.

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 17

Templates

- [Templates, on page 333](#)

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 43: 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 44: 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 335 .
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 336 .
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 336 .
Duplicate template	<p>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.</p> <p>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 336.</p>

Actions	Description
Delete template	<p>Allows you to delete a template. You can delete more than one template in a single instance.</p> <p>You can delete the user-defined templates. However, you cannot delete the predefined templates</p> <p>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.</p> <p>Note Select multiple templates to delete them at the same instance.</p> <p>To delete the template permanently, delete the template that is located in your local directory: <code>C:\Cisco Systems\dcn\ndfc\data\templates\</code>.</p>
Import	<p>Allows you to import a template from your local directory, one at a time. For more information, see Importing a Template, on page 338.</p>
Import as Zip	<p>Allows you to import .zip file, that contains more than one template that is bundled in a .zip format</p> <p>All the templates in the ZIP file are extracted and listed in the table as individual templates.</p> <p>For more information, see Importing a Template, on page 338.</p>
Export	<p>Allows you to export the template configuration to a local directory location. You can export only one template at a time.</p> <p>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.</p>

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 336](#).
- 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, N5600, 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.	<ul style="list-style-type: none"> • CLI • POAP <p>Note POAP option is not applicable for Cisco Nexus Dashboard Fabric Controller LAN Fabric deployment.</p> <ul style="list-style-type: none"> • POLICY • SHOW • PROFILE • FABRIC • ABSTRACT • REPORT 	Yes

Property Name	Description	Valid Values	Optional?
templateSubType	Specifies the sub type associated with the template.		

Property Name	Description	Valid Values	Optional?
		<ul style="list-style-type: none"> • CLI <ul style="list-style-type: none"> • N/A • POAP <ul style="list-style-type: none"> • N/A • VXLAN • FABRICPATH • VLAN • PMN Note POAP option is not applicable for Cisco Nexus Dashboard Fabric Controller LAN Fabric deployment. • POLICY <ul style="list-style-type: none"> • VLAN • INTERFACE_VLAN • INTERFACE_VPC • INTERFACE_ETHERNET • INTERFACE_BD • INTERFACE_CHANNEL • INTERFACE_FC • INTERFACE_MGMT • INTERFACE_COBACK • INTERFACE_NVE • INTERFACE_VFC • INTERFACE_CHANNEL • DEVICE • FEX 	

Property Name	Description	Valid Values	Optional?
		<ul style="list-style-type: none"> • NIRA_FABRIC_LINK • NIER_FABRIC_LINK • INTERFACE • SHOW <ul style="list-style-type: none"> • VLAN • INTERFACE_VLAN • INTERFACE_VPC • INTERFACE_ETHNET • INTERFACE_BD • INTERFACE_CHANNEL • INTERFACE_FC • INTERFACE_MGMT • INTERFACE_LOOPBACK • INTERFACE_NVE • INTERFACE_VFC • INTERFACE_VLAN • DEVICE • FEX • NIRA_FABRIC_LINK • NIER_FABRIC_LINK • INTERFACE • PROFILE <ul style="list-style-type: none"> • VXLAN • FABRIC <ul style="list-style-type: none"> • NA 	

Property Name	Description	Valid Values	Optional?
		<ul style="list-style-type: none"> • ABSTRACT <ul style="list-style-type: none"> • VLAN • INTERFACE_VLAN • INTERFACE_VPC • INTERFACE_ETHNET • INTERFACE_BD • INTERFACE_CHANNEL • INTERFACE_FC • INTERFACE_MGMT • INTERFACE_LOOPBACK • INTERFACE_NVE • INTERFACE_VFC • INTERFACE_PORTCHANNEL • DEVICE • FEX • NIRA_FABRIC_LINK • NIER_FABRIC_LINK • INTERFACE • REPORT <ul style="list-style-type: none"> • UPGRADE • GENERIC 	

Property Name	Description	Valid Values	Optional?
contentType		<ul style="list-style-type: none"> • CLI <ul style="list-style-type: none"> • TEMPLATE_CLI • POAP <ul style="list-style-type: none"> • TEMPLATE_CLI Note POAP option is not applicable for Cisco Nexus Dashboard Fabric Controller LAN Fabric deployment • POLICY <ul style="list-style-type: none"> • TEMPLATE_CLI • PYTHON • SHOW <ul style="list-style-type: none"> • TEMPLATE_CLI • PROFILE <ul style="list-style-type: none"> • TEMPLATE_CLI • PYTHON • FABRIC <ul style="list-style-type: none"> • PYTHON • ABSTRACT <ul style="list-style-type: none"> • TEMPLATE_CLI • PYTHON • REPORT <ul style="list-style-type: none"> • PYTHON 	Yes
implements	Used to implement the abstract template.	Text	Yes

Property Name	Description	Valid Values	Optional?
dependencies	Used to select the specific feature of a switch.	Text	Yes
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 “_” Discrete numbers separated by “,” Example: 1-10,15,18,20	Yes
interface	Format: <if type><slot>[/<sub slot>]/<port> 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	<p>You can have a list of IPv4, IPv6, or a combination of both types of addresses.</p> <p>Example 1: 172.22.31.97, 172.22.31.99, 172.22.31.105, 172.22.31.109</p> <p>Example 2: 2001:0cb8:85a3:0000:0000:8a2e:0370:7334, 2001:0cb8:85a3:0000:0000:8a2e:0370:7335, 2001:0cb8:85a3:1230:0000:8a2f:0370:7334</p> <p>Example 3: 172.22.31.97, 172.22.31.99, 2001:0cb8:85a3:0000:0000:8a2e:0370:7334, 172.22.31.254</p>	Yes
ipAddressWithoutPrefix	<p>Example: 192.168.1.1</p> <p>or</p> <p>Example: 1:2:3:4:5:6:7:8</p>	No
ipV4Address	IPv4 address	No
ipV4AddressWithSubnet	Example: 192.168.1.1/24	No
ipV6Address	IPv6 address	No
ipV6AddressWithPrefix	<p>Example: 1:2:3:4:5:6:7:8</p> <p>22</p>	No
ipV6AddressWithSubnet	IPv6 Address with Subnet	No
ISISNetAddress	<p>Example:</p> <p>49.0001.00a0.c96b.c490.00</p>	No
long	Example: 100	No
macAddress	14 or 17 character length MAC address format	No
string	<p>Free text, for example, used for the description of a variable</p> <p>Example:</p> <pre>string scheduledTime { regularExpr=^([01]\d 2[0-3]):([0-5]\d)\$; }</pre>	No

Variable Type	Valid Value	Iterative?
string[]	Example: {a,b,c,str1,str2}	Yes
struct	<p>Set of parameters that are bundled under a single variable.</p> <pre>struct <structure name declaration > { <parameter type> <parameter 1>; <parameter type> <parameter 2>; ... } [<structure_inst1>] [, <structure_inst2>] [, <structure_array_inst3 []>; struct interface_detail { string inf_name; string inf_description; ipAddress inf_host; enum duplex { validValues = auto, full, half; }; }myInterface, myInterfaceArray[];</pre>	<p>No</p> <p>Note If the struct variable is declared as an array, the variable is iterative.</p>
wwn (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. 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 Example: 75.56, -8.5	Yes	Yes	Yes	Yes	Yes							
floatRange	range of signed real numbers Example: 50.5 - 54.75	Yes	Yes	Yes	Yes	Yes							
integer	signed number Example: 50, -75	Yes	Yes		Yes	Yes							
integerRange	Range of signed numbers Example: 50-65	Yes	Yes		Yes	Yes							
interface	specific interface Example: Ethernet 5/10	Yes	Yes				Yes	Yes	Yes	Yes			
interfaceRange		Yes	Yes				Yes	Yes	Yes	Yes			
ipAddr	IP address in IPv4 or IPv6 format	Yes											

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
ipAddressList	<p>You can have a list of IPv4, IPv6, or a combination of both types of addresses.</p> <p>Example 1: 172.23.9, 172.3.9, 172.3.15, 172.3.10</p> <p>Example 2: 172.16.57, 172.16.57, 172.16.57</p> <p>Example 3: 172.3.9, 172.3.9, 172.16.57, 172.3.29</p> <p>Note Separate the addresses in the list using commas and not hyphens.</p>	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
ipAd	IPv4 or IPv6 Address (does not require prefix)												
ip4Ad	IPv4 address	Yes											
ip4Sub	IPv4 Address with Subnet	Yes											
ip6Ad	IPv6 address	Yes											
ip6Sub	IPv6 Address with prefix	Yes											
ip6Sub	IPv6 Address with Subnet	Yes											
ip6Ad	Example: 4008:5:0												
long	Example: 100	Yes			Yes	Yes							
macAd	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 Example for string Regular expression string { }	Yes									Yes	Yes	Yes
string[]	string literals that are separated by a comma (,) Example: {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 params that are bundled under a single variable. struct <structure name declaration> > { <parameter type> <parameter 1>; <parameter type> <parameter 2>; ... } <struct1> [, <struct2> [, <struct3> []>;												
wnn	WWN address												

Example: Meta Property Usage

```
##template variables

integer VLAN_ID {
min = 100;
max= 200;
};

string USER_NAME {
defaultValue = admin123;
minLength = 5;
};

struct interface_a{
```

```

string inf_name;
string inf_description;
ipAddress inf_host;
enum duplex {
    validValues = auto, full, half;
};
}myInterface;

##

```

Variable Annotation

You can configure the variable properties marking the variables using annotations.



Note Variable Annotations are available for POAP only. However, the annotations do not impact on the template type 'CLI'.

The following annotations can be used in the template variable section.

Annotation Key	Valid Values	Description
AutoPopulate	Text	Copies values from one field to another
DataDepend	Text	
Description	Text	Description of the field appearing in the window
DisplayName	Text Note 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 Note Use this annotation only for the ipAddress variable.
IsManagementIP	“true” or “false” 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”	

Annotation Key	Valid Values	Description
IsSourceSwitchName	“true” or “false”	
IsSwitchName	“true” or “false”	
IsRMID	“true” or “false”	
IsVPCDomainID	“true” or “false”	
IsVPCID	“true” or “false”	
IsVPCPeerLinkPort	“true” or “false”	
IsVPCPeerLinkPortChannel	“true” or “false”	
IsVPCPortChannel	“true” or “false”	
Password	Text	Validates the password field
PeerOneFEXID	“true” or “false”	
PeerTwoFEXID	“true” or “false”	
PeerOnePCID	“true” or “false”	
PeerTwoPCID	“true” or “false”	
PrimaryAssociation		
ReadOnly	“true” or “false”	Makes the field read-only
ReadOnlyOnEdit	“true” or “false”	
SecondaryAssociation	Text	
Section		
UsePool	“true” or “false”	
UseDNSReverseLookup		
Username	Text	Displays the username field on the window
Warning	Text	Provides text to override the Description annotation

Example: AutoPopulate Annotation

```
##template variables
string BGP_AS;
@(AutoPopulate="BGP_AS")
```

```

    string SITE_ID;
##

```

Example: DisplayName Annotation

```

##template variables
@(DisplayName="Host Name", Description = "Description of the host")
String hostname;
@(DisplayName="Host Address", Description = " test description" IsManagementIP=true)
ipAddress hostAddress;
##

```

Example: IsMandatory Annotation

```

##template variables
@(IsMandatory="ipv6!=null")
ipV4Address ipv4;
@(IsMandatory="ipv4!=null")
ipV6Address ipv6;
##

```

Example: IsMultiLineString Annotation

```

##template variables
@(IsMultiLineString=true)
string EXTRA_CONF_SPINE;
##

```

IsShow Annotation

```

##template variables
boolean isVlan;
@(IsShow="isVlan==true")
integer vlanNo;
##

##template variables
boolean enableScheduledBackup;
@(IsShow="enableScheduledBackup==true",Description="Server time")
string scheduledTime;
##
The condition "enableScheduledBackup==true" evaluates to true/false

##template variables
@(Enum="Manual,Back2BackOnly,ToExternalOnly,Both")
string VRF_LITE_AUTOCONFIG;
@(IsShow="VRF_LITE_AUTOCONFIG!=Manual", Description="Target Mask")
integer DCI_SUBNET_TARGET_MASK
##
The condition "VRF_LITE_AUTOCONFIG!=Manual" matches string comparison to evaluate to true
or false

```

Example: Warning Annotation

```

##template variables
@(Warning="This is a warning msg")
    string SITE_ID;
##

```

Templates Content

This section includes the configuration commands and any parameters that you want to include in the template. These commands can include the variables declared in the template variables section. During the command generation process the variable values are substituted appropriately in the template content.



Note You must specify the commands that you include as if you were entering them in the global configuration command mode on any device. You must consider the command mode when you include commands.

Template content is governed by the usage of variables.

- **Scalar variables:** does not take a range or array of values which cannot be used for iteration (In the variable types table those marked iterate-able as 'No'). Scalar variables must be defined inside the template content.

```
Syntax: $$<variable name>$$
Example: $$USER_NAME$$
```

- **Iterative variables:** used for block iteration. These loop variable must be accessed as shown below inside the iteration block.

```
Syntax:@<loop variable>
Example:
foreach val in $$INTEGER_RANGE_VALUE$$ {
@val
}
```

- **Scalar Structure Variable:** Structure member variables can be accessed inside the template content.

```
Syntax: $$<structure instance name>.<member variable name>$$
Example: $$myInterface.inf_name$$
```

- **Array Structure Variable:** Structure member variables can be accessed inside the template content.

```
Syntax: $$<structure instance name>.<member variable name>$$
Example: $$myInterface.inf_name$$
```

In addition to the template variables, you can use the conditional and iterative command generation using the following statements:

- **if-else if-else Statement:** makes a logical decision in inclusion/exclusion of set of configuration command based on the value assigned for the variable in it.

```
Syntax: if(<operand 1> <logical operator> <operand 2>){
command1 ..
command2..
..
}
else if (<operand 3> <logical operator> <operand 4> )
{
Command3 ..
Command4..
..
}
else
{
```

```

Command5 ..
Command6..
..
}
Example: if-else if-else statement
if($$USER_NAME$$ == 'admin'){
Interface2/10
no shut
}
else {
Interface2/10
shut
}

```

- **foreach Statement:** used for iterating a block of commands. The iteration is performed based on the assigned loop variable value.

```

Syntax:
foreach <loop index variable> in $$<loop variable>$$ {
@<loop index variable> ..
}
Example: foreach Statement
foreach ports in $$MY_INF_RANGE$${
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 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: $(\text{Usage Inside Range}/\text{Max Entries}) * 100$.

Click **View Details** to display a view of resources allocated or set in each resource pool. For example, the detailed section for a **SUBNET** has information about the resources that have been allocated within the subnet.

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.



CHAPTER 18

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.

- [Logs, on page 363](#)

Logs

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 19

Backup and Restore

You can take a back up immediately anytime. You can also configure a scheduler to 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.

The table displays all the Backups and Upgrades restored on Cisco Nexus Dashboard Fabric Controller. Click on the refresh icon to refresh the entries in the table.

This section includes the following:



Note If there are no scheduled back up jobs, **No Schedule set** is displayed.

- [Scheduler](#), on page 365
- [Restore](#), on page 366
- [Backup Now](#), on page 367

Scheduler

To restore 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 back up 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** to take backup of configuration data.

- Choose **Full** to backup of all previous version data to this application.

- 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 back up directory.
- Step 7** Enter the **Encryption Key** to the backup file.
- Step 8** In the **Run on days** field, select the check box to schedule the back up job on one or more days.
- Step 9** In the **Start at** field, use the time picker to schedule the back up at a particular time.
The time picker is a 12-hour clock.
- Step 10** Click **Schedule backup** to run the back up job as per schedule.
-

Restore

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** to restore only configuration data.
 - Choose **Full** to restore all previous version data to this application.
- 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.
 - 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.

Step 4 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** to take backup of configuration data.
- Choose **Full** to backup of all previous version data to this application.

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.
 - 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.
- 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.
 - e. Click **Backup**.

After the backup is complete, the backup file is saved in the remote directory.



CHAPTER 20

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 369](#)

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. 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 choose certificates with the `.cer/.crt/.pem` file extension separately.



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 needs to verify the NX-API certificates offered by the switches. In other words, the CA certificates are only 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.

Assigning Switches and Installing Certificates

These certificates are only consumed by Cisco Nexus Dashboard Fabric Controller, and not installed on switches.

Unlinking and Deleting Certificates

The CA certificates do not require to be unlinked, as they are never installed on switches.

CA certificates can still be deleted because one may need to bring new certificates for a given CA.

From the **Actions** drop-down list, select **Delete**. The certificate is deleted from Nexus Dashboard Fabric Controller.

Enabling NX-API Certificate Verification

The NX-API certificate verification can be 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 offered by the switches. If the verification fails, the NX-API calls will also fail.



Note

- Verification of the NX-API certificates can not 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.
 - Once an NX-API call fails for a given switch because of verification issues, the toggle button can be used 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.
-

