



About Fabric Overview for LAN
Operational Mode Setups, Release
12.2.2/12.2.3

Table of Contents

New and Changed Information	1
Fabric Overview	2
Overview	3
Hosts Card	3
Flows Card	4
Switches	5
Set Role Field	9
Guidelines and Limitations for changing discovery IP Address	10
Changing Discovery IP Address	11
Links	13
Creating Intra-Fabric Links	14
Creating Inter-Fabric Links	16
Protocol View	18
Policies	20
Adding a Policy	23
Creating a Policy Group	25
Advertising PIP on vPC	26
Custom Maintenance Mode Profile Policy	26
Creating and Deploying Custom Maintenance Mode Profile Policy	27
Deleting Custom Maintenance Mode Profile Policy	28
Event Analytics	29
Alarms	29
Cleared Alarms	29
Events	30
IPFM Events	31
Recent Tasks	32
VRFs	33
VRFs	33
Creating a VRF	36
VRF Attachments	42
Networks	47
Networks	47
Creating Network for Standalone Fabrics	52
Network Attachments	55
Private VLANs	59
Guidelines and Limitations for Private VLANs over VXLAN	59
Enabling PVLAN for a Fabric	59
Configuring an Interface as a PVLAN Port	60
Creating a Network for Primary and Secondary VLANs	61
Attaching a Primary Network	63
Attaching a Secondary Network	64

Explicit and Implicit Detach	65
History	66
Viewing Deployment History	66
Viewing Policy Change History	66
Resources	68
Allocating a Resource	68
Releasing a Resource	70
Hosts	72
Discovered Hosts Summary	72
Discovered Hosts	72
Host Policies	73
Deployment Status	79
Create Host Policy	80
Host Alias	81
Create Host Alias	83
Applied Host Polices	83
Flows	85
Flow Status	85
IPFM and Generic Multicast Flow Status	85
Multicast NAT Visualization	86
Flow Policies	92
Deployment Status	98
Creating a Flow Policy	99
Flow Alias	100
Creating Flow Alias	101
Static Flow	102
Creating a Static Flow	103
Metrics	104
Multicast NAT	109
Prerequisites	109
NAT Modes	109
Adding a NAT Mode	111
Deleting a NAT Mode	112
Recirc Mappings	112
Adding Recirc Mapping	115
NAT Rules	116
Adding NAT Rule	118
Deleting NAT Rule	119
RTP/EDI Flow Monitor	120
Active Flows	121
Packet Drop	122
Drop History	122
Global Config	123

Switch Global Config	123
Deployment History	126
Deployment Status	126
IPFM VRF	127
Deployment History	129
Deployment Status	129
VRF (Generic Multicast)	131
Virtual Infrastructure	132
Copyright	133

New and Changed Information

The following table provides an overview of the significant changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Release Version		Feature	Description
NDFC 12.2.3	release	Introduction of VRF Lite tab	When creating a VRF, the fields related to VRF Lite are now consolidated under a single VRF Lite tab. See Creating a VRF for more information.
NDFC 12.2.2	release	Support added for creating VXLAN EVPN fabrics with a PIMv6 Underlay and TRM IPv6	In previous releases of NDFC, NDFC supported an IPv6 underlay with ingress replication (IR). Beginning with the NDFC 12.2.2 release, NDFC added support for multicast replication. Previously NDFC supported a standalone VXLAN IPv4 fabric. Beginning with NDFC 12.2.2, NDFC supports creating a Multi-Site Domain (MSD) fabric with VXLAN IPv6. Prior to NDFC 12.2.2, NDFC supported Tenant Routed Multicast (TRM) IPv4. With NDFC 12.2.2, NDFC added support for TRM IPv6. A new tab, TRM, is added on the Create VRF page for enabling forwarding of multicast traffic for IPv4 or IPv6. This feature is available for the following fabric types: ▪ Data Center VXLAN EVPN fabric ▪ BGP (eBGP EVPN) fabric ▪ VXLAN EVPN Multi-Site fabric For more information, see the following topics: ▪ Creating a VRF > TRM ▪ Creating Inter-Fabric Links ▪ "Configuring VXLAN EVPN Fabrics with a PIMv6 Underlay and TRMv6" section in Data Center VXLAN EVPN

Fabric Overview

The **Actions** drop-down list at the fabric level allows you to perform these tasks. Note that some options are not available for certain fabric types.

Actions	Description
Edit Fabric	- To edit a fabric, choose Actions > Edit Fabric. - The Edit Fabric page appears. Enter the necessary updates and click Save.
Add Switches	For more information, see Add Switches: LAN .
Recalculate and Deploy	- To deploy configuration changes, choose Actions > Recalculate and Deploy. - A progress window appears and a confirmation message displays.
More	
Deployment Enable	- From Fabrics Overview, choose Actions on main tab, choose More > Deployment Enable. - A confirmation window appears, click OK.
Deployment Disable	- From the Fabric Overview page, choose Actions on the main tab and choose More > Deployment Disable. - A confirmation window appears, click OK.
Backup Fabric	Refer to the "Backup Fabric" section in Backup and Restore: LAN for more information.
Restore Fabric	Refer to the "Restore Fabric" section in Backup and Restore: LAN for more information.
VXLAN OAM	For more information, see the section "Configuring VXLAN OAM" in Understanding LAN Fabrics.  This feature appears in the Actions drop-down list only for the Data Center VXLAN EVPN fabric, eBGP VXLAN fabric, External Connectivity Network, and the Enhanced Classic LAN fabric technologies, which support VXLAN OAM.
Configure End Point Locator	The Endpoint Locator (EPL) feature allows real-time tracking of endpoints within a data center. For more information, see Endpoint Locator .

The **Fabric Overview** page contains tabs that allow you to 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 if you've configured IPFM fabric.

- Flows

This tab is displayed only if you've configured IPFM fabric.

- Reports

Hosts Card

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](#).

Flows Card

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](#).

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.



For all non-nexus device only MD5 protocol option is supported for SNMPv3 authentication.

The **Switches** tab has following information of every switch you discover in the fabric:

- **Switch:** 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.
- **Mode:** Specifies the current mode of the switch.
- **Config Status:** Specifies the configuration status. Status will be either **In-Sync** or **Out-of-sync**.
- **Oper Status:** Specifies the operation health status of the switch.
- **Discovery Status:** Specifies the discovery status of the switch. After the discovery of the device, the discovery status changes to **Ok** in green.
- **Model:** Specifies the switch model.
- **vPC Role:** Specifies the vPC role of the switch.
- **vPC Peer:** Specifies the vPC peer of the switch.
- **Software Version:** Specifies the software version of the switch.
- **Up Time:** Specifies the number of days and time since the switch is online.

The **Switches** tab has the following operations on the **Action** drop-down list.

Action Item	Description
Add switches	Click this icon to discover existing or new switches to the fabric. This option is also available in the Manage > Inventory > Switches page. Choose Actions > Add switches to discover and add switches to the fabric. Refer to the following sections for more information: ▪ "Adding Switches" section in BGP Fabric : Provides information on adding switches to easy fabrics. ▪ "Discovering New Switches" section in External Connectivity Network: Provides information on adding Cisco Nexus switches to external fabrics. ▪ "Adding Non-Nexus Devices to External Fabrics" section in External Connectivity Network: Provides 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. You can deploy configurations for multiple devices using the Deploy option.  This option grays out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric. In a VXLAN Multi-site 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 through the Discovery option: ▪ Update Credentials: Updates device credentials such as authentication protocol, username, and password. ▪ Rediscover: Initiates the switch discovery process by Nexus Dashboard Fabric Controller afresh. ▪ Change Discovery IP: Changes the discovery IP address for the switches. ▪ Update VRF: Enables auto discovery of the VRFs associated with the interface of the switch with a discovery IP address, when importing a switch.
Set Role	Assigns roles for one or more devices of the same device type. See Set Role Field for more information.

Action Item	Description
vPC Pairing	Choose a switch and click vPC Pairing to create, edit, or unpair a vPC pair. Use this option only when you choose a Cisco Nexus switch. For more information on how to create a vPC pair in external fabrics, see Creating a vPC Setup in the External Fabric.  NDFC 12 does not allow you to create vPC pairing on Spine, Border Spine, Border Gateway Spine, Super Spine, Border Super Spine, and Border Gateway Super Spine roles.
ToR/Access Pairing	Enables ToR to Leaf pairing for the switches.
vPC Overview	Enables you to configure a vPC pair for the selected switch.
More	The following operations are provided under More .
Change Mode	You can change the mode of a switch from Normal to Managed and vice versa. You can choose to save the settings and deploy immediately or schedule it for later.
Provision RMA	Allows you to replace a physical switch in a fabric when using Cisco Nexus Dashboard Fabric Controller Easy Fabric mode.
Change Serial Number	Allows you to change switch serial number if the switches are pre-provisioned. While pre-provisioning devices, you can provide dummy values for the Serial number of the switch. After configuring network for preprovisioned devices in the form of policies, or links, or interfaces, or VRFs, or networks, the dummy serial number can be changed with the required appropriate serial number. Before changing the serial number of the switches, choose Actions > Recalculate and deploy to save the latest data on switch.  Change of serial number is allowed only for Nexus 9000 Series switches.
Copy Run Start	Performs an on-demand copy running-configuration to startup-configuration operation for one or more switches.  This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.
Reload	Reloads the selected switch.  This option is grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.

Action Item	Description
Restore Switch	The information you restore on a switch is extracted from the fabric backups. This does not restore any fabric intents and other configurations applied using the fabric settings. Only intents on the switch are restored. Therefore, after you restore a switch, it might go out-of-sync because the fabric intents are not restored. Perform a fabric restore to restore the intents as well. You can restore only one switch at a time. You cannot restore a switch if the fabric where it is discovered is part of an VXLAN Multi-site fabric.
Export	Introduced in NDFC release 12.2.1. Exports switch inventory information to a .csv file. Check the box next to one or more switches to select those switches, then click Actions > Export to export the inventory information for the selected switches. This option is also available in the Manage > Inventory > Switches page. In that screen, choose Actions > Export to export switch inventory information to a .csv file.
Show Commands	Executes 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 executes 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.
Delete switches	Removes the selected 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.

Set Role Field

The **Set Role** field assigns roles for one or more devices of the same device type. 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.

The following are the device roles supported for non-Nexus devices:

- Spine
- Leaf
- Edge Router (Use this role for VRF-Lite).
- Core Router
- Super Spine
- ToR

You can change the switch role from an existing role to a supported role if there are no overlays on the switches. Click **Recalculate and Deploy** to generate the updated configuration. The following changes are allowed for a switch role:

- Leaf to Border
- Border to Leaf
- Leaf to Border Gateway
- Border Gateway to Leaf
- Border to Border Gateway
- Border Gateway to Border
- Spine to Border Spine
- Border Spine to Spine
- Spine to Border Gateway Spine
- Border Gateway Spine to Spine
- Border Spine to Border Gateway Spine
- Border Gateway Spine to Border Spine

You cannot change the switch role from any leaf role to any spine role or from any spine role to any leaf role.

In case the switch role is not changed according to the allowed changes mentioned above, the following error message is displayed after you deploy.

Switch[<serial-number>]: Role change from <switch-role> to <switch-role> is not permitted.

You can change the switch role to the role that was set earlier, or set a new role, and configure the fabric.

If you have not created any policy template instances before performing **Recalculate and Deploy**, and there are no overlays, you can change the role of a switch to any other required role.

If you change the switch role of a vPC switch that is part of a vPC pair, the following error appears when you perform **Recalculate and Deploy**:

Switches role should be the same for VPC pairing. peer1 <serial-number>: [<switch-role>],
peer2 <serial-number>: [<switch-role>]

To prevent this scenario, change the switch roles of both the switches in the vPC pair to the same role.

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:
 - Data Center VXLAN EVPN
 - BGP Fabric
 - External
 - Classic LAN
 - 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:

1. Choose **Manage > Fabrics**.
2. Click on fabric names to view the required switch.

The **Fabric summary** slide-in pane appears.

3. Click **Launch** icon to view **Fabric Overview** window.
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.

5. Click the check box next to the **Switch** column and select the switch.



You can change the IP address for individual switch and not for multiple switches.

6. Choose **Actions > Change Discovery IP** on the switches tab area.

The **Change Discovery IP** window appears.

Similarly, you can navigate from **Manage > Inventory > Switches** tab. Choose a required switch, click **Actions > Discovery > Change Discovery IP**.

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.
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**.



The PM records associated with the old switch IP address will be purged and new record collections take an hour to initiate after the changes.

Links

You can add links between border switches of different fabrics (inter-fabric links) or between switches in the same fabric (intra-fabric links). You can only create an inter-fabric connection (IFC) for a switch that is managed by Nexus Dashboard Fabric Controller.

There are scenarios where you might want to define links between switches before connecting them physically. The links could be inter-fabric or intra-fabric links. Doing so, you can express and represent your intent to add links. The links with intent are displayed in a different color till they are actually converted to functional links. Once you physically connect the links, they are displayed as connected.

Management links might show up in the fabric topology as red colored links. To remove such links, right-click the link and click **Delete Link**.

The Border Spine and Border Gateway Spine roles are added to switch roles for border switches.

You can create links between existing and pre-provisioned devices as well by selecting the pre-provisioned device as the destination device.

Starting from Cisco NDFC Release 12.1.2e, parameters MTU, SPEED, Source Interface Description, Destination Interface Description, Source Interface Freeform Config, and Destination Interface Freeform Config are added to the existing **int_pre_provision_intra_fabric_link** template. These parameters are preserved on subsequent **Recalculate & Deploy** after the device has completed bootstrap and POAP.

The following table describes the fields that appear on the **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: • Creating Inter-Fabric Links • Creating Intra-Fabric Links

Action Item	Description
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.  You cannot update existing links. The Import Links icon is disabled for an external fabric.
Export	Select Actions > Export to export the links in a CSV file. The following details of links are exported: link template, source fabric, destination fabric, source device, destination device, source switch name, destination switch name, source interface, destination interface, and nvPairs. The nvPairs field consists JSON object.

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:

1. From the **Actions** drop-down list, click **Create**.

The **Link Management - Create Link** page appears.

2. From the **Link Type** drop-down list, 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 the 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 field 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.



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.



The **Source IP** and **Destination IP** fields do not appear if you choose the `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 gets converted into a configuration, but will not be pushed into the switch. After **Save & Deploy**, it is reflected in the running configuration.

Disable BFD Echo on Source Interface and **Disable BFD Echo on Destination Interface** - Check the check box to disable BFD echo packets on the source and the 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, see [Enabling Freeform Configurations on Fabric Switches](#).

3. Click **Save** at the bottom right part of the page.

You can see that the IFC is created and displayed in the list of links.

4. On the **Fabric Overview Actions** drop-down list, select **Recalculate Config**.

The **Deploy Configuration** page 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. It lists the pending configurations on the switch. The **Side-by-Side Comparison** tab displays the running configuration and the expected configuration side-by-side.

Close the **Pending Config** page.

5. From the **Fabric Overview Actions** drop-down list, click **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 page. The **Links** screen displays 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.



In external fabrics, inter-fabric links support BGW, Border Leaf/Spine, and edge router switches. To create inter-fabric links, perform the following steps:

1. From the **Actions** drop-down list, select **Create**.

The **Link Management - Create Link** page appears.

2. From the **Link Type** drop-down box, choose **Inter-Fabric** since you are creating an IFC. The screen changes correspondingly.

The fields for inter-fabric link creation are as follows:

Field	Description
Link Type	Choose Inter-Fabric to create an inter-fabric connection between two fabrics, over their border switches.
Link Sub-Type	This field populates the IFC type. From the drop-down list, choose VRF_LITE, MULTISITE_UNDERLAY, or MULTISITE_OVERLAY. For information about VXLAN MPLS interconnection, see MPLS SR and LDP Handoff. For information about routed fabric interconnection, see the section "Creating Inter-Fabric Links Between a Routed Fabric and an External Fabric" in Managing BGP-Based Routed Fabrics.
Link Template	The link template is populated. The templates are autopopulated with corresponding prepackaged default templates that are based on your selection.  You can add, edit, or delete user-defined templates. See Templates for more details.

Field	Description
Source Fabric	This field is prepopulated with the source fabric name.
Destination Fabric	Choose the destination fabric from this drop-down box.
Source Device and Source Interface	Choose the source device and Ethernet interface that connects to the destination device.
Destination Device and Destination Interface	Choose the destination device and Ethernet interface that connects to the source device. Based on the selection of the source device and source interface, the destination information is autopopulated based on Cisco Discovery Protocol information, if available. There is an extra validation that is performed to ensure that the destination external device is indeed part of the destination fabric.

3. Navigate to the **General Parameters** tab.

Field	Description
Source BGP AS#	In this field, the AS number of the source fabric is autopopulated.
Source Address/Mask IP	In this field, enter the IPv4 address with a netmask of the source interface that connects to the destination device.
Destination Address IP	In this field, enter the IPv4 address of the destination interface.
Source Address/Mask IPv6	In this field, enter the IPv6 address with a netmask of the source interface.
Destination Address IPv6	In this field, enter the IPv6 address of the destination interface.
Destination BGP ASN Address	Specifies the BGP autonomous system number for the destination fabric.
BGP Maximum Paths	Specifies the maximum number of iBGP/eBGP paths. The valid value is between 1 and 64.
Routing TAG	Specifies the routing tag associated with the interface IP.
Link MTU	Specifies the interface MTU for both ends of the inter-fabric link.

4. 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.

5. On the **Fabric Overview > Actions** drop-down list, select **Recalculate Config**.

The **Deploy Configuration** page displays.

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.

6. From the **Fabric Overview > Actions** drop-down list, select **Deploy Config**.

The pending configurations are deployed.

After ensuring that the progress is 100% in all the rows, click **Close** at the bottom part of the page. The **Links** page 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 fabrics of an MSD, then you can see the link in the MSD topology too.

What's next:

When you enable the VRF-Lite function using the ToExternalOnly method or Multisite function over MSD fabric, IFCs are automatically created between the (VXLAN fabric) border/BGW device and connected (external fabric) edge router or core device. When you remove the ER/core/border/BGW device, the corresponding IFCs (link PTIs) to/from that switch are deleted on Nexus Dashboard Fabric Controller. Next, Nexus Dashboard Fabric Controller removes the corresponding IFC configurations, if any, from the remaining devices on the next **Save & Deploy** operation. Also, if you want to remove a device that has an IFC and overlay extensions over those IFCs, you should undeploy all the overlay extensions corresponding to those IFCs for switch delete to be possible.

1. To undeploy VRF extensions, select the VXLAN fabric and the extended VRFs, and undeploy the VRFs on the VRF deployment page.
2. To delete the IFCs, delete the IFCs from the **Links** tab.
3. Ensure that the fabric switch names are unique. If you deploy VRF extensions on switches with the same name, it leads to an erroneous configuration.

The new fabric is created, the fabric switches are discovered in Nexus Dashboard Fabric Controller, the underlay networks that are provisioned on those switches, and the configurations between Nexus Dashboard Fabric Controller and the switches are synced.

The remaining tasks are:

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

Protocol View

This tab displays the protocols for the links in the selected Fabric.

The following table describes the fields that appear on **Protocol View** tab.

Field	Description
Fabric Name	Specifies the name of the fabric.
Name	Specifies the name of the link.

Field	Description
Is Present	Specifies if the link is present.
Link Type	Specifies the type of link.
Link State	Specifies the state of link.
UpTime	Specifies the time duration from when the link was up.

Policies

Cisco Nexus Dashboard Fabric Controller manages configuration of devices using Policies. Nexus Dashboard Fabric Controller policies are way of grouping all the required CLIs and variables to achieve certain configuration on the devices. These policies can be defined either using CLI commands or Python scripts. Nexus Dashboard Fabric Controller generates the configuration for a device based on the policies attached to the device.

With Release 12.1.3, Nexus Dashboard Fabric Controller provides the ability to create policy groups which can be applied to multiple switches. Policy groups let you create policies that define specific switch parameters that are common to switches and apply them to multiple switches in a fabric.

To access the **Policies** page:

1. On the Cisco Nexus Dashboard Fabric Controller UI, choose **Manage > Fabrics** and double-click on a fabric.

The **Fabric Overview** page opens.

2. Go to the **Policies** tab

The following table describes the fields that appear on the **Policies** page.

Field	Description
Template	Specifies the name of the policy template.
Description	Specifies the description, if available.  From Cisco NDFC Release 12.1.1e, change of serial number for the switch is allowed, both old and new serial numbers can be viewed in this column.
Content Type	Specifies for the template content type. The supported content types are TEMPLATE_CLI, PYTHON and PYTHON_CLI.
Switch	Specifies the name of the switch the policy has been applied to. If you are configuring a policy group, this field provides a link that specifies the number of switches that are linked to the policy. Click the link to open the policy group details dialog box with details such as the number of switches that are linked to the policy, the IP address, the fabric name, serial number, and mark deleted state.
Entity Name	Specifies the switch or the interface name to which the policy has been applied to.
Entity Type	Specifies if the entity is a switch or an interface.
Source	Specifies the source.

Priority	Specifies the priority. During an Edit Membership operation to remove one or more switches from an existing policy group which uses Template_CLI content type, the Priority column displays the value Mixed indicating that the policy group has mixed priorities and mark deleted states. Whereas, when you edit <i>switch_freeform</i> policies of Content Type PYTHON (where multiple CLI policy templates are combined with a common source), after an edit operation the system removes occurrence of the switch from the source policy and displays the source and the child policies as different entries. The Mark Deleted value for these switches in a child policy indicates the value true and the Priority indicates a negative value. For a policy group, click on the link to view the group policy details of all the associated switches.
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 column displays <i>true</i> indicating that the policy is marked for deletion. All the configurations for a policy with the Mark Deleted value <i>true</i> will be negated. The Generated Config for the policy displays the configuration to be removed from the switch. For a policy group, click on the link to view the group policy details of all the associated switches.
Policy ID	Specifies the policy ID. The policy ID for a policy group begins with the term POLICY-GROUP. While searching for a policy group, you can filter the policy ID using this term.
IP Address	Specifies the IP address of the switch. If you are configuring a policy group, this field provides a link that specifies the IP addresses for the number of switches that are linked to the policy. Click on the link to view the group policy details of all the associated switches.
Serial Number	Specifies the serial number of the switch. If you are configuring a policy group, this field provides a link that specifies the serial numbers for the switches that are linked to the policy. Click on the link to view the group policy details of all the associated switches.
Created On	New field added in NDFC release 12.2.3. Specifies the date on which the policy was created.
Modified On	New field added in NDFC release 12.2.3. Specifies the date on which the policy was modified.

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

Action Item	Description
Add Policy	Allows you to create the following types of policies: • Regular policies. To add a regular policy, see Adding a Policy. • Policy Group. To add a policy group, see Creating a Policy Group.
Edit Policy	To modify the policy, choose a policy from the table and choose Edit Policy.  The policies in <i>italics</i> cannot be edited. The value under the Editable and Mark Deleted columns for these policies will indicate false. You cannot perform Edit Policy for policies 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.
Edit Membership	Lets you edit the membership configuration for a policy group. You can add or remove switches from a policy group using this option. If you remove switches from a policy group, the Switch column in the policy details table still displays the original number of switches in the policy. However the Mark Deleted column in the details view dialog box displays <i>true</i>. You cannot immediately edit a policy after an edit membership operation. The system displays an error message indicating to deploy the pending membership configuration changes to the switch before proceeding with any other configuration changes. If you choose to not deploy the membership configuration changes and would like to edit the policy, ensure you perform Preview and proceed with the edit policy operation. Preview operation removes any pending configuration changes from the system. You cannot perform Edit Membership for policies whose Mark Deleted value is set to <i>true</i>.

Delete Policy	To delete policies, choose the policies from the table and choose Delete Policy. The following are the points to consider while deleting group policies: • For TEMPLATE_CLI policies, removing a policy group removes all the child policies from the switch. • For Python policies which has a source and multiple child policies, removing a policy group removes the source policy template instance (PTI) from the switch and displays only the child policies. The system shows the Generated Config as negative for both the child policies. You cannot delete the child policies without deploying the configuration. The child policies are deleted automatically after deploying all the pending configuration.  A warning appears when you delete policies whose Mark Deleted values are set to <i>true</i>. Deleting a TEMPLATE_CLI policy removes the policy directly from the switch and sets the Mark Deleted value to <i>true</i>. When you delete policies whose Mark Deleted values are set to <i>true</i>, these entries are only removed from the NDFC database; the configs are not deployed to the switch. These policies do not have any intent and hence you need not deploy the config to the switch.
Generated Config	To view the delta of configuration changes made by every user, select policies from the table and choose Generated Config.
Push Config	To apply the policy configuration to the device, select policies from the table and choose Push Config. 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 apply configuration for a Python policy. You cannot perform a Push Config for policies whose Mark Deleted value is set to <i>true</i>.

Adding a Policy

To add a policy, perform the following steps:

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

The **Fabric Overview** window appears.

2. On the **Policies** tab for a fabric, choose **Actions > Add Policy**.

The **Create Policy** page appears.

3. Select the required switch and click **Next**.

You must deploy the switch in pending state.

4. Click **Choose Template** and choose the appropriate policy template and click **Select**.

From Cisco NDFC Release 12.1.2e, new templates **ipv4_prefix_list** and **ipv6_prefix_list** are added to the template list.

From Cisco NDFC Release 12.1.2e, you can enable or disable PTP high-correction notification when the system encounters a high-correction event. Whenever the correction value exceeds the configured value then that correction is called a high-correction. By default, a high-correction notification is disabled. Enable it manually to generate the notification. Perform the following steps to enable the high-correction notification:

- a. Check the **Enable PTP Telemetry** check box to enable telemetry for PTP.
- b. Check the **Is Large-Scale Fabric?** check box to generate the high-correction notification.

If there are more than 35 devices in a fabric, PTP events will be used if the switch version is 9.3(5) or higher, or else PTP correction data will be pushed periodically.

- c. Enter the wait time between two successive notifications in the **PTP High-Correction Interval** field.

The duration value is in seconds.

- d. Set correction range threshold value (ns) in the **PTP Correction Range** field.

The default is 100000 (100us).

5. Enter the priority value for the policy in the **Priority** field.

The applicable values are from 1 to 1000. The default value is 500. A lower number in the **Priority** field indicates that there is a higher priority for the generated configuration and POAP startup-configuration. For example, features are 50, route-maps are 100, and vpc-domain is 200.

6. Depending on the policy template that you have selected, enter all the necessary field values to create a policy and click **Save**.
7. If you have selected the **ipv4_prefix_list** or **ipv6_prefix_list**, perform the following steps to include the prefix-list entries:

- a. Enter the required name in the **Prefix List Name** field.
- b. On the **Prefix-list Entries** card, click **ActionsAdd**.

The **Add Item** window appears.

- c. Configure the mandatory fields on the **Add Item** dialog box and click **Save**.
- d. Repeat the step to add the required number of prefix-list entries.



The value in the **Sequence Number** must be higher than the previous prefix-list entry. If not, an error message is displayed.

- e. Select the appropriate prefix-list entry and click **Actions > Insert Above** to insert a new prefix-list entry.



The value in the **Sequence Number** must be lower than the below prefix-list entry. If not, an error message is displayed.

Creating a Policy Group

Policy groups provides a method of configuring and managing switches collectively. This feature enables you to create group policies for switches that share common configurations. You can create a policy group and add multiple switches to the policy at the time of creating the policy group or later. Similarly, Policy groups also let you edit or delete policies for multiple switches simultaneously.

To create a policy, perform the following steps:

1. Choose **LAN > Fabrics** and double-click on the required fabric.

The **Fabric Overview** window appears.

2. On the **Policies** tab for a fabric, choose **Actions > Add Policy**.

The **Create Policy** page appears.

3. To create a policy group, select the required switches to which you need to apply the policy and click **Next**.

Ensure you select switches that are part of the same fabric.

4. Enter the priority value for the policy in the **Priority** field.

The applicable values are from 1 to 1000. The default value is 500. A lower number in the **Priority** field indicates that there is a higher priority for the generated configuration and POAP startup-configuration. For example, the priority for vPC related policies are as follows: base_feature_vpc is 100, vpc-domain_mgmt is 150, for policies for interfaces on vPC (int_vpc_peer_link_po) is 202.

5. Use the toggle switch to enable or disable the **Group** option, as required.

If you have selected multiple switches, the **Group** toggle switch is enabled by default. If you select one switch initially and choose to add additional switches later, you can select the **Group** toggle switch to create a policy group and add additional switches later.

Not all templates provide support for creating policy groups. If you select a template that does not support policy group, the system generates an error message. Ensure you uncheck the **Group** toggle switch and create regular policies for templates that do not support policy group.

6. Click **Choose Template** and choose the appropriate policy template and click **Select**.

The available policy templates are TEMPLATE_CLI, PYTHON and PYTHON_CLI.

Note that nested Python policies are not supported. Additionally, when configuring policy groups, make sure you do not add policies that can be applied only on a single switch. Select policies that can be applied on multiple switches.

7. Depending on the policy template that you have selected, enter all the necessary field values to create a policy and click **Save**.

The new policy group appears in the **Fabric Overview > Policies** page.

8. To deploy the configuration to the switches, select the new policy that you have created and choose **Actions > Push Config**. Alternatively, to deploy the configuration, navigate to the **Fabric Overview** page, choose **Actions > Deploy** on the **Switches** tab.

Note that **Push Config** option does not go through configuration compliance check. It should be used only when you want to deploy commands which are ignored during configuration compliance checks.

Advertising PIP on vPC

Choose required fabric on LAN Fabric window and Navigate to **Edit Fabric > VPC**, check the **vPC advertise-pip** check box to enable the Advertise PIP feature on all vPCs in a fabric. Choose the **vpc_advertise_pip_jython** policy to enable Advertise PIP feature on specific vPCs in a fabric.

Note the following guidelines:

- If advertise-pip is not globally enabled or vPC peer is not using fabric peering, only then the vpc_advertise_pip_jython policy can be created on specific peers.
- The policy vpc_advertise_pip_jython can be applied only when switches are part of vPC pairing.
- Ensure that you configure **vpc advertise-pip** command during maintenance window as it involves BGP next-hop rewrite. Enabling this feature EVPN type 5 uses Switch Primary IP as next-hop while EVPN type 2 continue to use Secondary IP.
- Disabling advertise pip for a fabric doesn't affect this policy.
- Unpairing of switches deletes this policy.
- You can manually delete this policy from the peer switch where it was created.

To advertise PIP on vPC:

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

The **Fabric Overview** window appears.

2. On the **Fabric Overview** window, choose **Policies > Add Policy** and then select a switch with vPC.
3. Click **Actions > Add** and choose the switch from the **Switch List** drop-down list. Choose **vpc_advertise_pip_jython** policy template and enter the mandatory parameters data.



You can add this policy on one vPC peer, and it will create respective commands for vpc advertise on both peers.

4. Click **Save**, and then deploy this policy.

Custom Maintenance Mode Profile Policy

NDFC configures only a fixed set of BGP and OSPF isolate CLIs in the maintenance mode profile when you place a switch in maintenance mode. You can create a custom_maintenance_mode_profile policy with customized configurations for maintenance mode and normal mode profiles, deploy the

policy to the switch, and then move the switch to maintenance mode.

Creating and Deploying Custom Maintenance Mode Profile Policy

To create and deploy a custom maintenance mode profile policy from **Web UI > Switches**, perform the following procedure.

1. Select the desired switch and launch **Switch Overview**.
2. On the **Policies** tab, select **Actions > Add Policy** to add a new policy.
3. On the **Create Policy** page, click **Choose Template**.
4. Select **custom_maintenance_mode_profile** from the **Select Policy Template** list.
5. Fill in the **Maintenance mode profile contents** with the desired configuration CLIs.

Example:

```
configure maintenance profile maintenance-mode
ip pim isolate
```

Fill in the **Normal mode profile contents** with the desired configuration CLIs.

Example:

```
configure maintenance profile normal-mode
no ip pim isolate
configure terminal
```

6. Click **Save**.
7. From Switch Overview, click **Actions > Preview**.
8. Click on **Pending Config** lines to view the **Pending Config** and **Side-by-Side Comparison**.
9. Click **Close**.
10. From Switch Overview, click **Actions > Deploy**. Click **Deploy All** to deploy the new policy configuration on the switch.

Click **Close** after the deployment is complete.

11. Select the policy and select **Actions > More > Change Mode**.
12. In the **Mode** drop-down list, choose **Maintenance**.
13. Click **Save and Deploy Now** to move the switch to maintenance mode.

When you apply the default maintenance profile to a device functioning as an anycast border gateway (BGW) within a VXLAN fabric, it can result in dropping network traffic specifically for multi-fabric BUM traffic. This issue affects a subset of VNIs, particularly those for which the given BGW is designated as the forwarder. To address this issue, the maintenance mode profile must always use **include-local** for BGP isolation. In such cases, **custom_maintenance_mode_profile** policy must be created and deployed following the above steps, and then the necessary configuration CLIs should be modified.

The following is a sample `custom_maintenance_mode_profile` policy content for an anycast BGW in a VXLAN fabric with an OSPF underlay and multicast replication mode.

Normal mode profile contents example:

```
configure maintenance profile normal-mode
router ospf UNDERLAY
  no isolate
router bgp 65001
  no isolate include-local
no ip pim isolate
```

Maintenance mode profile contents example:

```
configure maintenance profile maintenance-mode
ip pim isolate
router bgp 65001
  isolate include-local
router ospf UNDERLAY
  isolate
```

Deleting Custom Maintenance Mode Profile Policy

The switch has to be moved to active/operational or normal mode before deleting the custom maintenance mode profile policy. To delete a custom maintenance mode profile policy from **Web UI > Switches**, perform the following procedure.

1. Select the desired switch and launch **Switch Overview**.
2. From **Switch Overview > Actions > More > Change Mode**.
3. In the Mode drop-down list, choose **Normal**.
4. Click **Save and Deploy Now** to move the switch to normal mode.
5. After the switch has been moved to normal mode, select the **custom_maintenance_mode_profile** policy that has to be deleted.
6. Choose **Actions > Edit Policy**.
7. Choose **Actions > Delete Policy** and click **Confirm** to mark the Policy for deletion.

The **Mark Deleted** column shows **true** indicating that the policy is marked for deletion.

8. Again, choose **Actions > Delete Policy** and click **Confirm** to delete the Policy.
9. From Switch Overview, click **Actions > Deploy**. Click **Deploy All** to delete the policy configuration on the switch.

Click **Close** after the deployment is complete.

Event Analytics

The event analytics section includes the following topics.

You can view the **Event Analytics** tab from the **Fabric Overview** dashboard as well as from the **Switch Overview** dashboard.

Alarms

1. Navigate to the **Manage > Fabrics** page.
2. Double-click on a fabric.

The **Fabric Overview** dashboard appears.

3. Click **Event Analytics > Alarms**.

The following table describes the fields that appear on the **Alarms** tab.

Field	Description
ID	(Optional) Specifies the ID of the alarm.
Severity	Specifies the severity of the alarm.
Source	Specifies the IP address of the alarm 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.

The following table describes the action items, in the **Actions** drop-down list, that appear on the **Fabric Overview > Event Analytics > Alarms** page.

Action Item	Description
Acknowledge	Select one or more events from the table and choose the 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 the Unacknowledge icon to acknowledge the event information for the fabric.
Clear	Select an event and choose Clear to clear the alarm.
Delete Alarm	Select an event and choose Delete Alarm to delete the alarm.

Cleared Alarms

1. Navigate to the **Manage > Fabrics** page.

2. Double-click on a fabric.

The **Fabric Overview** dashboard appears.

3. Click **Event Analytics > Cleared Alarms**.

The following table describes the fields that appear on the **Cleared Alarms** tab.

Field	Description
ID	(Optional) Specifies the ID of the alarm.
Severity	Specifies the severity of the alarm.
Source	Specifies the IP address of the alarm 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.
Cleared Time	Specifies the time at which the alarm was cleared.

The following table describes the action item, in the **Actions** drop-down list, that appears on the **Event Analytics > Cleared Alarms** tab.

Action Item	Description
Delete Alarm	Select an alarm and choose Delete Alarm to delete the alarm.

Events

1. Navigate to the **Manage > Fabrics** page.
2. Double-click on a fabric.

The **Fabric Overview** dashboard appears.

3. Click **Event Analytics > Events**.

The following table describes the fields that appear on the **Events** tab.

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 the syslog facility. The NDFC facility represents events generated by NDFC internal services and Simple Network Management Protocol (SNMP) traps generated by the switches. The syslog facility represents the machine process that created the syslog messages.

Field	Description
Type	Specifies how the switch or the fabric are managed.
Count	Specifies the number of times the event has occurred.
First Seen	Specifies the time when the event was created.
Last Seen	Specifies the time when the same event was seen 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** drop-down list, that appear on the **Fabric Overview > Event Analytics > Events** page.

Action Item	Description
Acknowledge	Select one or more events from the table and choose the 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 the Unacknowledge icon to acknowledge the event information for the fabric.
Delete Event	Select an event and choose Delete Event to delete the event.
Add Suppressor	Select an event and choose Add Suppressor to add a rule to the event. You can provide a name for the rule. Using the Scope options, you can add this rule to all the fabrics, or particular elements or all elements.
Event Setup	Allows you to create a new event.

IPFM Events



The **IPFM Events** tab displays for IPFM fabric templates only.

You can filter IPFM events by date or by the following attributes:

- **Switch**
- **Distinguished Name (DN)**
- **First Seen**
- **Last Seen**
- **Count**

You configure IPFM events from the **Fabric Controller > Admin > System Settings > Server Settings > IPFM** tab. NDFC maintains IPFM event history for seven days by default. For more information, see [Overview and Initial Setup of Cisco NDFC LAN](#).

1. Navigate to the **Manage > Fabrics** page.
2. Double-click on a fabric.

The **Fabric Overview** dashboard appears.

3. Click **Event Analytics > Events > IPFM Events**.

The following table describes the fields that appear on the **IPFM Events** tab.

Field	Description
Switch	Specifies the name of the switch.
Distinguished Name (DN)	Specifies the distinguished name (DN) for the IPFM event. Click on a Distinguished Name (DN) to view additional information about the DN or to filter information by Identifier , Reason , Time , or Status .
First Seen	Specifies the time when the event was first seen.
Last Seen	Specifies the time when the same event was seen last.
Count	Specifies the number of times the event has occurred.

Recent Tasks

On the **Recent Tasks** tab, you can view the changes made for the event analytics.



When the device is rebooted, the recent task details are erased.

1. Navigate to the **Manage > Fabrics** page.
2. Double-click on a fabric.

The **Fabric Overview** dashboard appears.

3. Click **Event Analytics > Recent Tasks**.

The following table describes the fields that appear on the **Recent Tasks** tab.

Field	Description
Fabric	Specifies the name of the fabric.
Task Name	Specifies the name of operation done on fabric recently.
Task Description	Specifies the description of task done on fabric.
Duration	Specifies the time duration of the task.
Completed/Progress	Specifies the progress details, whether the task is completed 100% or still in progress.

VRFs

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

1. Choose **Manage > Fabrics**.
2. Click on a fabric to open the **Fabric** slide-in pane.
3. Click the **Launch** icon.

The **Fabric Overview > Overview** page appears.

4. Click the **VRFs** tab.



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, see section "Overlay Mode" in [Understanding LAN Fabrics](#).

You can view the VRF details in the **VRFs** tab and VRF attachment details in the **VRF Attachments** tab.

This section contains the following:

VRFs

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

1. Choose **Manage > Fabrics**.
2. Click on a fabric to open the **Fabric** slide-in pane.
3. Click the **Launch** icon.

The **Fabric Overview** page appears.

4. Click on the **VRFs** tab.

Use the **VRFs** tab to create, edit, delete, attach, detach, import, export, and deploy configurations for VRFs. You can create networks only after creating a VRF except when you use Layer 2 to create a network.

VRFs 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.

5. Click the table header to sort the entries in alphabetical order for the selected parameter.

The following table describes the action items, in the **Actions** drop-down list, that appear on the **VRFs** tab on the **Fabric Overview** page.

VRFs Actions and Description

Action Item	Description
Create	Allows you to create a new VRF. For more information, see Creating a VRF .
Edit	Allows you to edit the selected VRF. 1. To edit a VRF, check the check box next to the VRF that you want to edit and choose Edit. On the Edit VRF page, you can edit the parameters. 2. Click Save to retain the changes or click Close to discard the changes.
Multi-Attach	Allows you to attach multiple switches to a VRF. 1. To attach a VRF, check the check box next to the VRF that you want to attach the switches to and choose Multi-Attach. On the Multi-Attach of VRFs page, you can specify the switches that you want to attach to the VRF. 2. Click Next to proceed to the next step in the wizard or click Cancel to discard the changes. The Summary page displays with the Proceed to Full Switch Deploy (Recommended) button selected. 3. Click Save. The Deploy Configuration page appears. 4. Click Deploy All. The Deploy Configuration page appears with an updated status of SUCCESS, Status Description, and Progress indicator. 5. Click Close. The attached VRF displays as DEPLOYED in the VRF Status column.

Action Item	Description
Multi-Detach	Allows you to detach selected switches from a VRF. 1. To detach a VRF, check the check box next to the VRF that you want to detach and choose Multi-Detach. On the Multi-Detach of VRFs page, you can specify the switches that you want to detach from the VRF. 2. Check the check box for the switch that you want to detach from the VRF. 3. Click Next. The Summary page displays with the Proceed to Full Switch Deploy (Recommended) button selected. 4. Click Save. The Deploy Configuration page appears with the selected switch. 5. Click Deploy All. Click Close.
Deploy	Allows you to deploy the configuration for the selected VRF. 1. To deploy a VRF configuration, check the check box next to the VRF for which you want to deploy the configuration and choose Deploy. On the Deploy Configuration page, you can deploy the specified VRF configuration. 2. Click Deploy or click Close to discard the changes.
Import	Allows you to import VRF information exported to a .csv file for the fabric. 1. To import VRF information exported to a .csv file, choose Import. The Import VRFs dialog box appears. 2. Browse to the directory and select the .csv file that contains the VRF information. 3. Click OK. The VRF information is imported and displayed on the Fabric Overview > VRFs page.

Action Item	Description
Export	Allows you to export VRF information to a .csv file. The exported .csv file contains information pertaining to each VRF, including the configuration details that you saved during the creation of the VRF. To export VRF information, choose Export. The VRF .csv file is exported to your local directory. The file name is appended with the date and time at which the file was exported.  You can use the exported .csv file for reference or use it as a template for creating new VRFs.
Delete	Allows you to delete a selected VRF. You can select multiple VRF entries and delete them at the same time. 1. To delete a VRF, check the check box next to the VRF that you want to delete and choose Delete. A warning message appears asking whether you want to delete the VRF(s). 2. Click Confirm to delete or click Cancel to retain the VRF. A message appears that the selected VRFs are deleted successfully.

Creating a VRF

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

1. Choose **Manage > Fabrics**.
2. Double-click on the fabric to open the **Fabric Overview** page.
3. Click on the **VRFs** tab.
4. On the **VRFs** tab, click **Actions > Create**.

The **Create VRF** page appears.

5. On the **Create VRF** page, enter the required details in the mandatory fields. The available fields vary based on the fabric type.

The fields on the **Create VRF** page are:

Field	Description
VRF Name	Specifies a VRF name automatically or allows you to enter a name. The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:). For MSD fabrics, the values for VRF or network are the same for the fabric.
VRF ID	Specifies the ID for the VRF or enter an ID for the VRF.

Field	Description
VLAN ID	Specifies the corresponding tenant VLAN ID for the network or enter an ID for the VLAN. If you want to propose a new VLAN for the network, click Propose VLAN .
Default Security Action	Available starting in NDFC release 12.2.2, related to the security groups feature introduced in this release. For more information on security groups, see Configuring Security for VXLAN EVPN Fabrics. The following options are available for the Default Security Action field: ▪ Unenforced: Default setting. There are no default security policies in place and therefore no security action is taken on the traffic that passes. ▪ Enforced Permit: Based on a permit list model, where traffic will be permitted on this VRF by default. You can configure granular contracts to deny specific traffic. ▪ Enforced Deny: Based on a deny list model, where traffic will be denied on this VRF by default. You can configure granular contracts to permit specific traffic.
Default Security Tag for VRF	Available starting in NDFC release 12.2.2, related to the security groups feature introduced in this release. For more information on security groups, see Configuring Security for VXLAN EVPN Fabrics. The value in this field will be automatically populated from the Security Tag Pool. This tag is used by default for traffic on this VRF unless the IP address or VLAN of that traffic is specifically classified as a selector under a security group.
VRF Template	A default universal template is auto-populated. This is applicable for leaf switches only.
VRF Extension Template	A default universal extension template is auto-populated. This allows you to extend this network to another fabric. The methods are VRF Lite, Multi Site, and so on. The template is applicable for border leaf switches and BGWs.

6. Enter the necessary field values or edit pre-filled fields, as required.

The tabs and their fields on the page are explained in the following sections.

- [General Parameters](#)
- [Advanced](#)
- [TRM](#)
- [Route Target](#)
- [VRF Lite](#) (NDFC 12.2.3 release only)

7. Click **Create** to create the VRF or click **Close** 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 shown as **NA** because the VRF is created but is not yet deployed. Double-click on the configured VRF to bring up the **VRF Overview** information.



If you did not associate a VLAN with a VRF when you created the VRF using these instructions, you will see **NA** displayed in the **VRF Overview** for the VRF, even if a VLAN was associated with the VRF through another process (for example, if you disabled the **Enable L3VNI w/o VLAN** setting).

Now that the VRF is created, you can create and deploy networks on the devices in the fabric.

General Parameters

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

Field	Description
VRF VLAN Name	Enter the VLAN name for the VRF.
VRF Interface Description	Enter a description for the VRF interface.
VRF Description	Enter a description for the VRF.

Advanced

Field	Description
VRF Interface MTU	Specifies the VRF interface MTU.
Loopback Routing Tag	If a VLAN is associated with multiple subnets, then this tag is associated with the IP prefix of each subnet. Note that this routing tag is associated with overlay network creation also.
Redistribute Direct Route Map	Specifies the redistribute direct route map name.
Max BGP Paths	Specifies the maximum number of BGP paths. The valid value is between 1 and 64.
Max iBGP Paths	Specifies the maximum number of iBGP paths. The valid value is between 1 and 64.
Enable IPv6 link-local Option	Select the check box to enable the IPv6 link-local option under the VRF SVI. If this check box is unchecked, IPv6 forwarding is enabled.

Field			Description
Enable L3VNI w/o VLAN			Beginning with NDFC release 12.2.1, check the box to enable the L3VNI w/o VLAN configuration. The default value of this field comes from the fabric-level field Enable L3VNI w/o VLAN. The default setting for this field varies depending on the following factors: • For existing VRFs, the default setting is disabled (the Enable L3VNI w/o VLAN box is unchecked). • For newly-created VRFs, the default setting is inherited from the fabric settings. • This field is a per-VXLAN fabric variable. For VRFs that are created from a VXLAN EVPN Multi-Site fabric, the value of this field will be inherited from the fabric setting in the child fabric. You can edit the VRF in the child fabric to change the value, if desired.
			Beginning with NDFC 12.2.3, the following fields have moved to the VRF Lite tab.
Advertise Host Routes			Check this check box to control advertisement of /32 and /128 routes to edge routers.
Advertise Route	Default		Check this check box to control advertisement of default route internally. To allow inter-subnet communication between end hosts in different VXLAN fabrics, where the subnets are present in both fabrics, you must disable the Advertise Default Route feature (clear the Advertise Default Route check box) for the associated VRF. This will result in /32 routes for hosts in both fabrics. For example, Host1 (VNI 30000, VRF 50001) in Fabric1 can send traffic to Host2 (VNI 30001, VRF 50001) in Fabric2 only if the host route is present in both fabrics. When a subnet is present in one fabric only then the default route is sufficient for inter-subnet communication.
Config Static 0/0 Route			Check this check box to control configuration of static default route.
BGP Password	Neighbor		Specifies the VRF-Lite BGP neighbor password.
BGP Password Encryption Type	Key		From the drop-down list, select the encryption type.
Enable Netflow			Allows you to enable netflow monitoring on the VRF-Lite sub-interface. Note that this is supported only if netflow is enabled on the fabric.
Netflow Monitor			Specifies the monitor for the VRF-Lite netflow configuration. To enable netflow on a VRF-Lite sub-interface, you must enable netflow at the VRF level and VRF extension level. Check the Enable_IFC_Netflow check box in the VRF attachment while you edit an extension to enable netflow monitoring. For more information, see the section "Netflow Support" in Understanding LAN Fabrics.

TRM

Beginning with NDFC 12.2.2, the **TRM** tab was added for configuring Tenant Routed Multicast (TRM) IPv6. The existing IPv4 TRM fields were moved from the **Advanced** tab to the **TRM** tab.

For more information on configuring TRM IPv6, see the section "Configuring VXLAN Fabrics with an IPv6 Multicast Underlay and TRM with IPv6" in [Data Center VXLAN EVPN](#).

For more information on TRM, see [Configuring Tenant Routed Multicast](#).

Field	Description
IPv4 TRM Enable	Check the check box to enable IPv4 TRM. If you enable IPv4 TRM, and provide the RP address, you must enter the underlay multicast address in the Underlay Mcast Address field.
NO RP	Check the check box to disable RP fields. You must enable IPv4 TRM to edit this check box. If you enable No RP, then the Is RP External, RP Address, RP Loopback ID, and Overlay Mcast Groups fields are disabled.
Is RP External	Check this check box if the RP is external to the fabric. If this check box is not checked, RP is distributed in every VTEP.
RP Address	Specifies the IP address of the RP.
RP Loopback ID	Specifies the loopback ID of the RP, if Is RP External is not enabled.
Underlay Multicast Address	Specifies the multicast address associated with the VRF. The multicast address is used for transporting multicast traffic in the fabric underlay.  The multicast address in the Default MDT Address for TRM VRFs field on the fabric settings page is auto-populated in this field. You can override this field if a different multicast group address should be used for this VRF.
Overlay Mcast Groups	Specifies the multicast group subnet for the specified RP. The value is the group range in the <code>ip pim rp-address</code> command. If the field is empty, 224.0.0.0/24 is used as the default.
TRMv6 Enable	Check this check box to enable IPv6 TRM.
TRMv6 No RP	Check this check box to disable RP fields in TRMv6 as only PIM-SSM is used.
Is TRMv6 RP External	Check this check box if the RP is external to the fabric in TRMv6.
TRMv6 RP Address	Enter the IPv6 address of the TRMv6 RP multicast traffic.
Overlay IPv6 Mcast Groups	Specifies the IPv6 multicast group subnet for the specified TRMv6 RP. The value is the group range in the <code>ipv6 pim rp-address</code> command. If the field is empty, ff00::/8 is used as the default.
Enable MVPN inter-as	Check this check box to use the inter-AS keyword for the Multicast VPN (MVPN) address family routes to cross the BGP autonomous system (AS) boundary. This option is applicable if you enabled the TRM option.

Field	Description
Enable IPv4/IPv6 TRM BGW MSite	Check this check box to enable IPv4 or IPv6 TRM on BGW multisite.

Route Target

Field	Description
Disable RT Auto-Generate	Check this check box to disable RT auto-generate for IPv4, IPv6 VPN/EVPN/MVPN.
Import	Specifies one VPN route target or a comma-separated list of VPN route targets to import.
Export	Specifies one VPN route target or a comma-separated list of VPN route targets to export.
Import EVPN	Specifies one EVPN route target or a comma-separated list of EVPN route targets to import.
Export EVPN	Specifies one EVPN route target or a comma-separated list of EVPN route targets to export.
Import MVPN	Specifies one MVPN route target or a comma-separated list of MVPN route targets to import.
Export MVPN	Specifies on MVPN route target or a comma-separated list of MVPN route targets to export.  By default, Import MVPN and Export MVPN fields are disabled. Check the IPv4 TRM Enable or the TRMv6 Enable check box to enable these fields.

VRF Lite

This tab is introduced in NDFC release 12.2.3.

Field	Description
Advertise Host Routes	Check this check box to control advertisement of /32 and /128 routes to edge routers.
Advertise Default Route	Check this check box to control advertisement of default route internally. To allow inter-subnet communication between end hosts in different VXLAN fabrics, where the subnets are present in both fabrics, you must disable the Advertise Default Route feature (clear the Advertise Default Route check box) for the associated VRF. This will result in /32 routes for hosts in both fabrics. For example, Host1 (VNI 30000, VRF 50001) in Fabric1 can send traffic to Host2 (VNI 30001, VRF 50001) in Fabric2 only if the host route is present in both fabrics. When a subnet is present in one fabric only then the default route is sufficient for inter-subnet communication.
Config Static 0/0 Route	Check this check box to control configuration of static default route.

Field	Description
BGP Neighbor Password	New field introduced in NDFC 12.2.3. Specifies the VRF-Lite BGP neighbor password.
BGP Password Key Encryption Type	New field introduced in NDFC 12.2.3. From the drop-down list, select the encryption type.
BGP Allow AS In	New field introduced in NDFC 12.2.3. Check the box to accept the AS-path, even if it contains ASN configured on this border switch.
BGP Allow AS In Number	New field introduced in NDFC 12.2.3. Specifies the number of occurrences of ASN allowed in the AS-path. (Min:1, Max:10, Default:3).
BGP AS Override	New field introduced in NDFC 12.2.3. Check this box to override matching ASN while sending a BGP update.
BGP Disable Peer AS Check	New field introduced in NDFC 12.2.3. Check this box to disable checking of peer ASN while advertising route to that BGP peer. NX-OS specific.
BGP Allow Inbound Soft Reconfiguration Always	New field introduced in NDFC 12.2.3. Check this box to store raw BGP updates locally to reprocess policies without route refresh.
Enable Netflow	Allows you to enable netflow monitoring on the VRF-Lite sub-interface. Note that this is supported only if netflow is enabled on the fabric.
Netflow Monitor	Specifies the monitor for the VRF-Lite netflow configuration. To enable netflow on a VRF-Lite sub-interface, you must enable netflow at the VRF level and VRF extension level. Check the Enable_IFC_Netflow check box in the VRF attachment while you edit an extension to enable netflow monitoring. For more information, see the section "Netflow Support" in Understanding LAN Fabrics.

VRF Attachments

UI Navigation

The following options are applicable only for switch fabrics, VXLAN EVPN fabrics, and VXLAN EVPN Multi-Site fabrics.

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-out pane. Click the **Launch** icon. Choose **Fabric Overview > VRFs > VRF Attachments**.
- Choose **Manage > 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.

VRF Attachments Table Fields and Description

Field	Description
VRF Name	Specifies the name of the VRF.

VRF ID	Specifies the ID of the VRF.
VLAN ID	Specifies the VLAN ID.
Switch	Specifies the name of the switch.
Status	Specifies the status of VRF attachments, for example, pending, NA, deployed, out-of-sync, and so on.
Attachment	Specifies whether the VRF attachment is attached or detached.
Switch Role	Specifies the switch role. For example, for the fabric created using the Campus VXLAN EVPN fabric template, the switch role is specified as either leaf, spine, or border.
Fabric Name	Specifies the name of the fabric to which the VRF is attached or detached.
Loopback ID	Specifies the loopback ID.
Loopback IPV4 Address	Specifies the loopback IPv4 address.
Loopback IPV6 Address	Specifies the loopback IPv6 address.  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.

VRF Attachments Actions and Description

Action Item	Description
History	Allows you to view the deployment and policy change history of the selected VRF. 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. 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. To view the history of a VRF attachment, check the check box next to the VRF name and select History. The History window appears. Click the Deployment History*or *Policy Change History tabs as required. You can also click the Detailed History link in the Commands column of the Deployment History tab to view the command execution details (comprising configuration, status, and CLI response) for the host.
Edit	Allows you to view or edit the VRF attachment parameters such as interfaces that you want to attach to the selected VRF. To edit the VRF attachment information, check the check box next to the VRF name that you want to edit. Select Edit. In the Edit VRF Attachment*window, edit the required values, attach or detach the VRF attachment. Click the *Edit link to edit the CLI freeform config for the switch, and click Save to apply the changes or click Cancel to discard the changes. The edited VRF attachment is shown in the table on the VRF Attachments horizontal tab of the VRFs tab in the Fabric Overview window.

Action Item	Description
Preview	Allows you to preview the configuration of the VRF attachments for the selected VRF.  This action is not allowed for attachments that are in deployed or NA status. To preview the VRF, check the check box next to the VRF name and choose Preview from Actions drop-down list. The Preview Configuration window for the fabric appears. You can preview the VRF attachment details such as the VRF name, fabric name, switch name, serial number, IP address, and role, VRF status, pending configuration, and progress of the configuration. Click the lines link in the Pending Config column to view the lines for which the configuration is pending. Click Close.
Deploy	Allows you to deploy the pending configuration of the VRF attachments, for example, interfaces, for the selected VRF.  This action is not allowed for attachments that are in deployed or NA status. To deploy a VRF, check the check box next to the VRF name and choose Deploy from Actions drop-down list. The Deploy Configuration window for the fabric appears. You can view the details such as the VRF name, fabric name, switch name, serial number, IP address, and role, VRF status, pending configuration, and progress of the configuration. Click the lines link in the Pending Config column to view the lines for which the configuration is pending. Click the Deploy button. The status and progress of the deployment is displayed in the VRF Status and Progress columns. After the deployment is completed successfully, close the window.

Action Item	Description
Import	Allows you to import information about VRF attachments for the selected fabric. To import the VRF attachments information, choose Import. Browse the directory and select the .csv file that contains the VRF attachments information. Click Open and then click OK. The VRF information is imported and displayed in the VRF Attachments horizontal tab on the VRFs tab in the Fabric Overview window.
Export	Allows you to export the information about VRF attachments to a .csv file. The exported file contains information pertaining to each VRF, including the fabric it belongs to, whether the LAN is attached, the associated VLAN, serial number, interfaces, and freeform configuration details that you saved for VRF attachments. To export VRF attachments information, choose the Export action. Select a location on your local system directory to store the VRF information and click Save. The VRF information file is exported to your local directory. The file name is appended with the date and time at which the file was exported.
Quick Attach	Allows you to immediately attach an attachment to the selected VRF. You can select multiple entries and attach them to a VRF at the same instance. To quickly attach any attachment to a VRF, choose Quick Attach from Actions drop-down list. A message appears to inform that the attach action was successful.
Quick Detach	Allows you to detach the selected VRF immediately from an attachment, for example, a fabric. You can select multiple entries and detach them from an attachment at the same instance. To attach any attachment to a VRF quickly, choose Quick Detach from Actions drop-down list. A message appears to inform that the detach action was successful.

Networks

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

1. Choose **Manage > Fabrics**.
2. Click on a fabric to open the **Fabric** slide-in pane.
3. Click the **Launch** icon.

The **Fabric Overview** page displays.

4. Choose the **Networks** tab.



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](#).

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, see the section "Overlay Mode" in [Understanding LAN Fabrics](#).

For more information on creating interface groups and attaching networks, see the section "Interface Groups" in [Add Interfaces: LAN](#).

You can view the network details in the **Networks** tab and network attachment details in the **Network Attachments** tab.

Networks

The following table describes the action items, in the **Actions** drop-down list, that appear on the **Networks** page.

Networks Actions and Description

Action Item	Description
Create	Allows you to create a new network for the fabric. For instructions about creating a new network, see Creating Network for Standalone Fabrics .
Edit	Allows you to view or edit the selected network parameters. 1. To edit the network information, check the check box next to the network name that you want to edit and choose Edit. 2. On the Edit Network page, edit the required values and click Save to apply the changes or click Close to discard the changes. The edited network is shown in the table in the Networks tab of the Fabric Overview page.

Action Item	Description
Multi-Attach	Allows you to attach multiple switches and interfaces to the network at the same time. 1. To attach the selected switches and interfaces to the network, select the check box next to the network name that you want to attach and choose Multi-Attach. 2. On the Multi-Attach of Networks page, check the check boxes for the switches that you want to attach to the network and click Next. 3. In the Select Interfaces area, check the check boxes for the interfaces that you want to attach to the network and click Next. The Summary page displays with the Proceed to Full Switch Deploy (Recommended) option selected. 4. Click Save. The Deploy Configuration page appears with the selected switch. 5. Click Deploy All. The Deploy Configuration page appears with an updated status of SUCCESS, Status Description, and Progress indicator. 6. Click Close. The attached network displays as DEPLOYED in the Network Status column in the Networks tab of the Fabric Overview page.

Action Item	Description
Multi-Detach	Allows you to detach multiple switches from the network at the same time. 1. To detach the selected switches from the network, select the check box next to the network name that you want to detach and choose Multi-Detach. 2. On the Multi-Detach of Networks page, check the check boxes for the switches that you want to detach from the network and click Next. The Summary page displays with the Proceed to Full Switch Deploy (Recommended) option selected. 3. Click Save. The Deploy Configuration page appears with the selected switch. 4. Click Deploy All. The Deploy Configuration page appears with an updated status of SUCCESS, Status Description, and Progress indicator. 5. Click Close. The detached network displays as NA (not attached) in the Network Status column in the Networks tab of the Fabric Overview page.
Deploy	Allows you to deploy the pending configuration for associating the switches or interfaces to the network. 1. To deploy a network, check the check box next to the network name that you want to deploy and choose Deploy. The Deploy Configuration page for the fabric appears. You can view the details such as the network name, fabric name, switch name, serial number, IP address, and role, network status, pending configuration, and progress of the configuration. 2. Click the Lines link in the Pending Config column to view the lines of the pending configuration. The Pending Config dialog box appears. 3. Click Cancel after you have viewed the pending configuration. 4. On the Deploy Configuration page, click the Deploy button. The status and progress of the deployment displays in the Network Status and the Progress columns. 5. After the deployment completes successfully, close the page.

Action Item	Description
Import	Allows you to import network information for the fabric. 1. To import network information, choose Import. The Import Networks dialog box appears. 2. Browse to the directory with the .csv file that contains the host IP address and corresponding unique network information. 3. Click OK. The host aliases are imported and displayed in the Networks tab of the Fabric Overview page.
Export	Allows you to export network information to a .csv file. The exported file contains information pertaining to each network, including the fabric it belongs to, the associated VRF, the network templates used to create the network, and all other configuration details that you saved during network creation. You can use the exported .csv file for reference or use it as a template for creating new networks. 1. To export network information, choose Export. 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.  The Networks tab displays network names based on the number of rows per page. You can view network names based on the options in the Rows per page drop-down list. When you use the Export option, NDFC exports the network names as displayed per page. If you have a large number of network names and you want to export all of your network names, you need to navigate to each page and export each page individually. 2. Before importing the file, update new records in the .csv file. 3. Ensure that the networkTemplateConfig field contains the JSON Object.
Delete	Allows you to delete the network. You can select multiple network entries and delete them at the same time. 1. To delete a network for the fabric, select the check box next to the network name that you want to delete and choose Delete. A Warning dialog box appears. 2. Click Confirm to delete the network.

Action Item	Description
Add to Interface Group	Allows you to add the network to an interface group. You can select multiple network entries and add them to an interface group at the same time. 1. To add the selected networks to the interface group that you want, check the check box for the network name you want and click Add to Interface Group. 2. On the Add to Interface Group page, click the networks link and verify whether the selected networks are present on the Selected Networks page and then click Cancel. 3. Either choose an Interface Group from the drop-down list or click Create Interface Group. 4. On the Create Interface Group page, provide the interface group name, select the interface type, and then click Create to save the changes or click Close to close the page and discard the changes. 5. On the Add to Interface Group page, click Save to save the changes or click Close to close the page and discard the changes. The interface group displays in the Interface Group column in the Networks tab of the Fabric Overview page.
Remove from Interface Group	Allows you to remove the network from an interface group. You can select multiple network entries and remove them from an interface group at the same time. 1. To remove the selected networks from the interface group, check the check box for the network name for which you want to remove the interface group. 2. Choose Remove from Interface Group. 3. On the Remove from Interface Group page, click the networks link and verify whether the selected networks are present on the Selected Networks page and then click Cancel. 4. On the Remove from Interface Group page, click Remove to remove the networks from the interface group or click Close to discard the changes. The interface group no longer displays in the Interface Group column in the Networks tab of the Fabric Overview page.

Networks Table Fields and Description

Field	Description
Network Name	Specifies the name of the network.
Network ID	Specifies the Layer 2 VNI of the network.
VRF Name	Specifies the name of the Virtual Routing and Forwarding (VRF).
IPv4 Gateway/Prefix	Specifies the IPv4 address with a subnet.

Field	Description
IPv6 Gateway/Prefix	Specifies the IPv6 address with a subnet.
Network Status	Displays the status of the network.
VLAN ID	Specifies the VLAN id.
Interface Group	Specifies the interface group.

Creating Network for Standalone Fabrics

Before creating networks, ensure that you have created a VRF for the fabric. However, if you have chosen Layer 2 on the **Create Network** page, then you do not require a VRF. For more information, see [VRFs](#).

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

1. On the **Networks** tab, click **Actions > Create**.

The **Create Network** page appears.

2. On the **Create Network** page, enter the required details in the mandatory fields. The available fields vary based on the fabric type.

The fields on the **Create Network** page are:

Field	Description
Network ID and Network Name	Specifies the Layer 2 VNI and the name of the network. The network name should not contain any white spaces or special characters, except underscore (_) and hyphen (-). The corresponding Layer 3 VNI (or VRF VNI) is generated along with VRF creation.
Layer 2 Only	Specifies whether the network is Layer 2 only.
VRF Name	Allows you to select the Virtual Routing and Forwarding (VRF) from the drop-down list. If you want to create a new VRF, click Create VRF . The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).
VLAN ID	Specifies the corresponding tenant VLAN ID for the network. If you want to propose a new VLAN for the network, click Propose VLAN .
Network Template	A default universal template is auto-populated. This is only applicable for leaf switches.
Network Extension Template	A default universal extension template is auto-populated. This allows you to extend this network to another fabric. The methods are VRF Lite, Multi Site, and so on. The template is applicable for border leaf switches and BGWs.
Generate Multicast IP	Click to generate a new multicast group address and override the default value.

The tabs and their fields in the screen are explained in the following sections.

- [General Parameters](#)
- [Advanced](#)

The fields on the **General Parameters** tab are:

Field	Description
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.  If the network is a non-Layer 2 network, then it is mandatory to provide the gateway IP address.  If the same IP address is configured in the IPv4 Gateway and IPv4 Secondary GW1 or GW2 fields of the network template, Nexus Dashboard Fabric Controller does not show an error, and you will be able to save this configuration. However, after the network configuration is pushed to the switch, it would result in a failure as the configuration is not allowed by the switch.
IPv6 Gateway/Prefix List	Specifies the IPv6 address with subnet.
Vlan Name	Enter the VLAN name.
Interface Description	Specifies the description for the interface. This interface is a switch virtual interface (SVI).
MTU for L3 interface	Enter the MTU for Layer 3 interfaces range 68 - 9216.
IPv4 Secondary GW1	Enter the gateway IP address for the additional subnet.
IPv4 Secondary GW2	Enter the gateway IP address for the additional subnet.
IPv4 Secondary GW3	Enter the gateway IP address for the additional subnet.
IPv4 Secondary GW4	Enter the gateway IP address for the additional subnet.

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

Field	Description
ARP Suppression	Select the check box to enable the ARP Suppression function.

Field	Description
Ingress Replication	The check box is selected if the replication mode is ingress replication.  Ingress replication is a read-only option on the Advanced tab. Changing the fabric setting updates the field.
Multicast-group-address	The multicast IP address for the network is autopopulated. Multicast group address is a per fabric instance variable. The number of underlay multicast groups supported is 128. If all networks are deployed on all switches, you need not use a different multicast group per L2 VNI or a network. Therefore, multicast group for all networks in a fabric remains same. Starting from Cisco NDFC Release 12.1.2e, a maximum of 16 DHCP relay servers for overlay networks are supported. Perform the following steps to include the DHCP relay server information: - On the DHCP Relay Server Information field, click Actions > Add. The ADD Item page appears. - Enter the Server IP V4 Address and Server VRF details and click Save. - Repeat the above steps to add the required number of DHCP relay server information.  When you upgrade to NDFC Release 12.1.2e and newer, the existing DHCP server configurations in the network definitions using the shipping overlay templates will be automatically updated to the new structure without any configuration loss.
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.
IPv4 TRM enable	Check the check box to enable TRM with IPv4. For more information, see Configuring Tenant Routed Multicast.
IPv6 TRM enable	Check the check box to enable TRM with IPv6. For more information, see Configuring Tenant Routed Multicast.
L2 VNI Route-Target Both Enable	Check the check box to enable automatic importing and exporting of route targets for all L2 virtual networks.

Field	Description
Enable Netflow	Enables netflow monitoring on the network. This is supported only if netflow is already enabled on fabric.
Interface Vlan Netflow Monitor	Specifies the netflow monitor specified for Layer 3 record for the VLAN interface. This is applicable only if Is Layer 2 Record is not enabled in the Netflow Record for the fabric.
Vlan Netflow Monitor	Specifies the monitor name defined in the fabric setting for Layer 3 Netflow Record .
Enable L3 Gateway on Border	Check the check box to enable a Layer 3 gateway on the border switches.

4. Click **Create**.

A message appears indicating that the network is created.

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

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

Network Attachments

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

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

Use this window to attach fabrics and interfaces to a network.

Network Attachments Table Fields and Description

Field	Description
Network Name	Specifies the name of the network.
Network ID	Specifies the Layer 2 VNI of the network.
VLAN ID	Specifies the VLAN ID.
Switch	Specifies the name of the switch.
Ports	Specifies the ports for the interfaces.
Status	Specifies the status of the network attachments, for example, pending, NA, and so on.
Attachment	Specifies whether the network attachment is attached or detached.
Switch Role	Specifies the switch role. For example, for the fabric created using the Campus VXLAN EVPN fabric template, the switch role is specified as either leaf, spine, or border.

Fabric Name	Specifies the name of the fabric to which the network is attached or detached.
-------------	--

The following table describes the action items, in the **Actions** drop-down list, that appears in the **Network Attachments** horizontal tab on the **Networks** tab in the **Fabric Overview** window.

Network Attachments Actions and Description

Action Item	Description
History	Allows you to view the deployment and policy change history of the selected network. 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. 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. To view the history of a network attachment, select the check box next to the network name and choose the History action. The History window appears. Click the Deployment History or Policy Change History tabs as required. Click the Detailed History link in the Commands column of the Deployment History tab to view the command execution details (comprising configuration, status, and CLI response) for the host.

Edit	Allows you to view or edit the network attachment parameters such as interfaces that you want to attach to the selected network. To edit the network attachment information: 1. Check the check box next to the network name that you want to edit and choose Actions > Edit. 2. In the Edit Network Attachment window: - o Attach or detach the network attachment - o Click the Edit link to edit the CLI freeform config for the switch - o Edit the required values  Beginning with NDFC release 12.2.3, when editing a network attachment in a VXLAN fabric on a switch with border or border gateway shown in the Switch Role column, you can also provide comma-separated route targets to import or export at the switch level in the Import Route Target and Export Route Target fields. This is supported for both Nexus 9000 and Catalyst 9000 switches. Then click Save to apply the changes or click Cancel to discard the changes. The edited network attachment is shown in the table on the Network Attachments horizontal tab of the Networks tab in the Fabric Overview window.
Preview	Allows you to preview the configuration of the network attachments for the selected network.  This action is not allowed for attachments that are in deployed or NA status. To preview the network, check the check box next to the network name and choose Preview from Actions drop-down list. The Preview Configuration window for the fabric appears. You can preview the network attachment details such as the network name, fabric name, switch name, serial number, IP address, and role, network status, pending configuration, and progress of the configuration. Click the lines link in the Pending Config column to view the lines for which the configuration is pending. Click Close.

Deploy	Allows you to deploy the pending configuration of the network attachments, for example, interfaces, for the selected network.  This action is not allowed for attachments that are in deployed or NA status. To deploy a network, check the check box next to the network name and choose Deploy from Actions drop-down list. The Deploy Configuration window for the fabric appears. You can view the details such as the network name, fabric name, switch name, serial number, IP address, and role, network status, pending configuration, and progress of the configuration. Click the lines link in the Pending Config column to view the lines for which the configuration is pending. Click the Deploy button. The status and progress of the deployment is displayed in the Network Status and Progress columns. After the deployment is completed successfully, close the window.
Import	Allows you to import information about network attachments for the selected fabric. To import the network attachments information, choose Import. Browse the directory and select the .csv file that contains the network attachments information. Click Open and then click OK. The network information is imported and displayed in the Network Attachments horizontal tab on the Networks tab in the Fabric Overview window.
Export	Allows you to export the information about network attachments to a .csv file. The exported file contains information pertaining to each network, including the fabric it belongs to, whether the LAN is attached, the associated VLAN, serial number, interfaces, and freeform configuration details that you saved for network attachments. To export network attachments information, choose the Export action. Select a location on your local system directory to store the network information and click Save. The network information file is exported to your local directory. The file name is appended with the date and time at which the file was exported.
Quick Attach	Allows you to immediately attach an attachment to the selected network. You can select multiple entries and attach them to a network at the same instance.  Interfaces cannot be attached to a network using this action. To quickly attach any attachment to a network, choose Quick Attach from Actions drop-down list. A message appears to inform that the attach action was successful.

Quick Detach	Allows you to immediately detach the selected network from an attachment, for example, a fabric. You can select multiple entries and detach them from an attachment at the same instance. To quickly detach any attachment to a network, choose Quick Detach from Actions drop-down list. A message appears to inform that the detach action was successful. After quick detach, the switch status is not computed when there is no deploy. Post deploy, the configuration compliance calls at entity level (interface or overlay).
--------------	--

Private VLANs

Starting with Release 12.1.3, Cisco Nexus Dashboard Fabric Controller introduces support for Private VLAN (PVLAN) over VXLAN.

A Private Virtual Local Area Network (PVLAN) is a VLAN that isolates a Layer 2 port from the other ports in the same broadcast domain or subnet. PVLAN restricts Layer 2 traffic within a broadcast domain by segmenting the broadcast domain into multiple subdomains. A subdomain contains a PVLAN pair which includes a primary VLAN and one or more secondary VLANs. A PVLAN domain can have multiple PVLAN pairs, one for each subdomain. All VLAN pairs in a PVLAN domain share the same primary VLAN. A PVLAN domain can have only one primary VLAN.

Secondary VLANs provide Layer 2 isolation between ports within the same PVLAN. Although PVLANS provide host isolation at Layer 2, hosts can communicate with each other at Layer 3.

Guidelines and Limitations for Private VLANs over VXLAN

- This feature is supported with Data Center VXLAN EVPN, BGP, and External fabrics.
- This feature is supported over physical interface, port-channel interface and virtual port channel (vPC) interfaces on the switches in a fabric.
- This feature is supported on Layer 2 ToR interfaces.
- This feature provides support for **cli** and **config-profile** overlay modes for VRF and network configuration.
- This feature is supported only on VTEPs, and not supported on spine and super spine switches.
- This feature is not supported with the VXLAN EVPN Multi-site fabrics.
- This feature is not supported on Brownfield deployments.
- This feature is not supported on Interface groups on PVLAN interface.

Enabling PVLAN for a Fabric

1. In Cisco Nexus Dashboard Fabric Controller, choose **Lan > Fabrics**.
2. Choose **Actions > Create Fabric** and select the required template and click **Select**.

To enable PVLAN on an existing fabric, select the fabric name and choose **Actions > Edit Fabric**.

3. Go to the **Advanced** tab and check the **Enable Private VLAN (PVLAN)** checkbox.

Ensure that you have checked the **Enable EVPN VXLAN Overlay** check box in the **EVPN** tab of the BGP fabric. You can enable the **Enable Private VLAN (PVLAN)** checkbox only if you have enabled VXLAN EVPN mode in your fabric.

4. From the **PVLAN Secondary Network Template** list, select PVLAN template for the secondary network. The default is **Pvlan_Secondary_Network**.
5. Click **Save**.

A warning message appears prompting you to perform a Recalculate and Deploy.

6. Click **OK**.
7. Double-click the fabric to open the **Fabric Overview** window.
8. Choose **Actions > Recalculate and Deploy**.
9. Review the configurations after the **Config Preview** and click **Deploy All**.

Performing a relcalculate and deploy enables **feature private-vlan** command on all the VTEPs and TORs.



You cannot disable PVLAN feature in a fabric, if there are any PVLAN networks or PVLAN interface policies configured.

Configuring an Interface as a PVLAN Port

Before You Begin

Ensure that you have enabled PVLAN feature for the fabric.

Perform the following steps to configure a PVLAN port:

1. In Cisco Nexus Dashboard Fabric Controller, choose **Lan > Fabrics**.
2. Double-click the fabric name to open the **Fabric Overview** page.
3. On the **Interfaces** tab, do one of the following:
 - For an Ethernet interface, select the required interface and choose **Actions > Edit**.
 - For a Port Channel or virtual Port Channel (vPC) interface, choose **Actions > Create Interface**.
4. Under the **Policy** field, click the policy link to select the required PVLAN interface policy.
5. In the **Select Attached Policy Template** dialog box, choose the required interface policy template and click **Select**.

The following are the supported PVLAN interface policies:

- **int_pvlan_host**: Specifies the interface template for creating a PVLAN port on an Ethernet interface.
- **int_port_channel_pvlan_host**: Specifies the interface template for creating a PVLAN port-channel interface.
- **int_vpc_pvlan_host**: Specifies the interface template for creating a vPC port for the PVLAN on a vPC pair.

After attaching the PVLAN policy to an interface, the **PVLAN** tab appears.

6. Configure all the necessary fields in the **PVLAN** tab.

The fields in the **PVLAN** tab are described in the following table.

Field	Description
PVLAN Mode	Specifies the PVLAN port type. The following are the supported types: • promiscuous • trunk promiscuous • host • trunk secondary
PVLAN Allowed Vlans	Configures a list of allowed VLANs on a PVLAN trunk port.
Native Vlan	Configures a VLAN to transport the untagged packets on PVLAN trunk ports. If there is no native VLAN configured, all untagged packets are dropped.
PVLAN Mapping	Displays the mapping between the primary VLAN and the secondary VLANs. The fields in this area are enabled only if you select promiscuous or trunk promiscuous as the PVLAN mode. You can configure multiple VLAN pairs for the PVLAN. To add new primary-secondary VLAN pair, choose Actions > Add .
PVLAN Association	Configures the association between the primary VLAN and the associated secondary VLANs. The fields in this area are enabled only if you select host or trunk secondary as the PVLAN mode. You can configure multiple VLAN pairs for the PVLAN. To add a new primary-secondary VLAN pair, choose Actions > Add .

7. When you have entered all the necessary information in the configuration fields, click **Save**.

An error message appears if you have not enabled PVLAN for the fabric. See [Enabling PVLAN for a Fabric](#) for the steps to enable PVLAN for the fabric.

For external fabrics, Cisco Nexus Dashboard Fabric Controller provides support for PVLAN only at the interface level. Before configuring PVLAN interfaces, if **feature private-vlan** is not already enabled on the switch, ensure that you add a PVLAN policy for the switch using the **feature-pvlan** policy template. Perform a **Recalculate and Deploy** and then follow the steps mentioned in this section to create PVLAN interfaces.

Creating a Network for Primary and Secondary VLANs

1. In Cisco Nexus Dashboard Fabric Controller, choose **Manage > Fabrics**.
2. From the list of available fabrics, double-click the PVLAN-enabled fabric.

The **Fabric Overview** page appears.

3. Navigate to the **Networks** tab and choose **Actions > Create**.

The **Create Network** window appears.

4. Enter the required details in the following fields. Some of the fields are auto-populated with default values. You can make changes, as required.

The fields in the **Create Network** window are:

Field	Description
Network Type	Click the Private (PVLAN) radio button. This radio button is available only if you have enabled private VLAN feature for the selected fabric.
Private Network Type	Specifies the VLAN type. Select one of the following options: ▪ Primary - Select the option to configure your network as the primary VLAN. You can configure only one primary VLAN in a PVLAN. ▪ Community - Select the option to configure a secondary VLAN to enable the hosts to communicate with each other as well as forward traffic to ports in the Primary VLAN. ▪ Isolated - Select the option to configure an isolated secondary VLAN that enables the hosts to only forward traffic to the ports in the primary VLAN.
Network Name	Specifies the name of the network. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-).
Layer 2 Only	Enables you to create a Layer 2 only network. This field is applicable only for primary VLANs.
Primary Network Name	Choose the name of the primary network from the list of configured primary networks. This field is applicable only when you are configuring a secondary VLAN.

Field	Description
VRF Name	Allows you to select the VRF that you have created for the fabric. When no VRF is created, this field appears as 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 (:). This field is applicable only for primary VLANs.
Network ID	Specifies the layer 2 Virtual Network Identifier (VNI) of the network.
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	Auto-populates the universal template for primary networks. For secondary networks, select the Pvlan_Secondary_Network template. This is only applicable for leaf switches.
Network Extension Template	Auto-populates the universal extension template for primary networks. For secondary networks, select the Pvlan_Secondary_Network template. This allows you to extend the network to another fabric. The VRF Lite extension is supported. The template is applicable for border leaf switches.

5. When you have entered all the necessary information in the configuration fields, click **Create**.

The table in the **Networks** tab displays all the newly created PVLAN networks.

What's next: Once you have primary and secondary networks configured, you can attach the networks to the switches.

Attaching a Primary Network

After creating the primary and secondary networks, you can attach the networks to the switches and their PVLAN interfaces. You can attach a primary network either explicitly or implicitly.

- **Explicit Attach/Detach** - Defines the method of manually attaching/detaching a network.
- **Implicit Attach/Detach** - Defines the method in which a network is attached/detached automatically because one of the members in a PVLAN primary-secondary pair undergoes an explicit attachment/detachment.

This section is optional for VTEPs that only have PVLAN host or trunk secondary ports. If you want to perform an implicit attach for your primary network, you can skip the section and proceed to

Attaching a Secondary Network.

Perform the following steps to attach a primary network explicitly.

1. In Cisco Nexus Dashboard Fabric Controller, choose **Manage > Fabrics**.
2. From the list of available fabrics, double-click the PVLAN-enabled fabric.

The **Fabric Overview** page appears.

3. Navigate to the **Networks** tab and double-click the primary network to open the **Network Overview** page.
4. On the **Network Attachments** tab, select the required networks and choose **Actions > Edit**.

The **Edit Network Attachment** page opens.

The table under **Available Interfaces for this device** displays all the promiscuous ports and the promiscuous trunk ports available in the device. Note that for a primary PVLAN network, only the promiscuous ports and the promiscuous trunk ports are displayed.

If you have ToR switches connected to the device, the pvlan interfaces on TOR switch will be displayed. If you select any TOR interface, the system adds PVLAN configuration to the TOR switch.

5. Use the toggle button to enable **Attach** and then click **Save**.
6. On the **Networks** tab, select the network and choose **Actions > Deploy**.

Attaching a Secondary Network

Perform the following steps to explicitly attach a secondary network.

1. In Cisco Nexus Dashboard Fabric Controller, choose **Manage > Fabrics**.
2. From the list of available fabrics, double-click the PVLAN-enabled fabric.

The **Fabric Overview** page appears.

3. Navigate to the **Networks** tab and double-click the secondary network to open the **Network Overview** page.
4. On the **Network Attachments** tab, select the required networks and choose **Actions > Edit**.

The **Edit Network Attachment** page opens.

The table under **Available Interfaces for this device** displays all the ports of type host and trunk secondary. For a secondary PVLAN network, only the host ports and the trunk secondary ports are displayed. Both community and isolated PVLAN ports are labeled as PVLAN host ports. A PVLAN host port is either a community PVLAN port or an isolated PVLAN port depending on the type of secondary VLAN with which it is associated.

If you have ToR switches connected to the device, the pvlan interfaces on TOR switch will be displayed. If you select any TOR interface, the system adds PVLAN configuration to the TOR switch.

Note that you cannot attach an interface group to a secondary network.

5. Use the toggle button to enable **Attach**, and then click **Save**.

If you have not already performed an **Attach** for your primary network, the system automatically attaches the primary network along with the secondary network. You can view the network status for both the primary and the secondary networks in the **Networks** tab of the **Fabric Overview** window.

When a secondary network is attached to a switch, it implicitly attaches to the other switches where its primary network is in explicit attach state, if the secondary network is not already attached.

6. On the **Networks** tab, select the network and choose **Actions > Deploy**.

Explicit and Implicit Detach

The steps to detach a network are similar to the steps for attaching a network. The following points describes how implicit and explicit detach feature works.

- When you detach a primary network in explicit state, the following occurs:
 - If there is no secondary network in explicit state on the switch, the primary network is detached along with all the associated secondary networks
 - If there is any secondary network in explicit state, the primary network does not detach but changes to implicit state
- When you detach a secondary network explicitly, the primary network detaches automatically (implicitly) if the following conditions are met:
 - If the primary network is in implicit attached state
 - If the detached secondary is the only secondary network for this primary network on this switch
 - If no other switch in the fabric has this secondary in explicit attach state, this secondary network also gets detached from the other switches

History

The history tab displays information about the deployment and policy change history. Choose **Manage > Fabrics**. Double-click a fabric name to open the **Fabric Overview** window and then click the **History** tab.

Viewing Deployment History

Deployment HistoryDeployment 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.



In certain scenarios after deploying configurations on a switch, the **Status** column displays the following error indicating a failure. However, the **Config Status** in **Manage > Inventory > Switches** page displays **In-Sync**. The reason for a conflicting status is due to inter-node connectivity issue that causes the keep alive between the sim-agent and the sim-master to fail resulting in an agent expiry event. Whereas the sim-agent is alive and deploying the configuration on the switch. This does not require any further action.

Deployment Failed to execute job for this device with " Reason: dcnm-sim-agent-xx expired.

Viewing Policy Change History

Policy Change HistoryDifferent 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.
Serial Number	Specifies the serial number.
Content Type	Specifies the content type.
User	Specifies the user.
Source	Specifies the source.

Resources

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:

1. Choose **Manage > Fabrics**.
2. Double-click a fabric name.

The **Fabric Overview** window appears.

3. Click the **Resources** tab.

4. Click **Actions > Allocate Resource** to allocate the resource.

The **Allocate Resource** window appears.

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.

6. Enter the entity name in the **Entity Name** field.

The embedded help gives example names for different scope types.

7. Enter the ID, IP address, or the subnet in the **Resource** field based on what pool type you chose in *Step 3*.

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:

1. Choose **Manage > Fabrics**.
2. Double-click a fabric name.

The **Fabric Overview** window appears.

3. Click the **Resources** tab.
4. Choose a resource that you want to delete.



You can delete multiple resources at the same time by choosing multiple resources.

5. Click **Actions > Release Resource(s)** to release the resource.

A confirmation dialog box appears.

6. Click **Confirm** to release the resource.

Hosts



This tab is only available on IPFM fabric when you have deployed IPFM on Nexus Dashboard Fabric Controller.

- Choose **Manage > 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 [Overview](#).

The **Hosts** tab includes the following tabs:

Discovered Hosts Summary

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Hosts > Discovered Hosts Summary**.
- Choose **Manage > 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.

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

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Hosts > Discovered Hosts**.
- Choose **Manage > 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.

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: ▪ 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

1. Navigate to the **Manage > Fabrics** page.
2. Click on a fabric name to open the **Fabric** slide-in pane.
3. Click the **Launch** icon.
4. Click **Fabric Overview > Hosts > Host Policies**.
5. Click **Manage > Fabrics**.
6. Double-click on a fabric name to open **Fabric Overview > Hosts > Host Policies**.

You can add policies to the host devices.

7. Navigate to **Host Policies** to configure the host policies.



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 the **Fabric Controller > Admin > System 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 on the **Host Policies** page.

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.



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 on the **Host Policies** page.



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 on this page 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 appear on the **Host Policies** page.

Host Policies Actions and Descriptions

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 .

Edit Host Policy	Allows you to view or edit the selected host policy parameters. To edit the host policy, select the check box next to the host policy that you want to delete and choose Edit Host Policy. On the Edit Host Policy page, 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 on the Host Policies page.  The changes made to host policy are applied immediately. If the policy is already applied to any device, the changes may impact the existing flows.
Delete Host Policy	Allows you to delete user-defined host policies.  ▪ Undeploy policies from all the switches before deleting them from Nexus Dashboard Fabric Controller. ▪ A default policy can be undeployed from the switches on which it is deployed. However, a custom policy can be deleted and undeployed. ▪ When you undeploy the default policies, all default policies are reset to have the default permission of Allow. 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. A delete host policy successful message appears at the bottom of the page.

Purge	Allows you to delete all custom policies without selecting any policy check box.  ▪ 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.
Import	Allows you to import host policies from a .csv file to Nexus Dashboard Fabric Controller.  After import, all policies imported from a .csv file are applied to all managed switches automatically. To import a host policies, choose Import. Browse the directory and select the .csv file that contains the host policy configuration information. The policy will not be imported if the format in the .csv file is incorrect. Click Open. The imported policies are automatically deployed to all the switches in the fabric.
Export	Allows you to export host policies from Nexus Dashboard Fabric Controller to a .csv file. To export host policies, choose Export. Select a location on your local system directory to store the host policy details file. Click Save. The host policy file is exported to your local directory. The filename is appended with the date on which the file is exported. The format of the exported file is .csv.
Deploy Selected Policies	Select this option to deploy only the selected policies to the switch.
Deploy All Custom Policies	Select this option to deploy all the custom or user-defined policies to the switch in a single instance. If the policies are deployed when the switch is rebooting, the deployment fails and a failed status message appears.
Deploy All Default Policies	Select this option to deploy all default policies to the switch.

Undeploy Selected Policies	Select this option to undeploy the selected policies. Select one or more check boxes next to the policy name. Select this option from the drop-down list to undeploy the selected policies.
Undeploy All Custom Policies	Select this option to undeploy all the custom or user-defined policies in a single instance.
Undeploy All Default Policies	Select this option to undeploy the default policies.
Redo All Failed Policies	The deployment of policies may fail due to various reasons. Select this option to deploy or undeploy all failed policies. 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 one policy from the drop-down list. Select this option to view the deployment history of the selected policy in the Deployment History pane. The policy name is shown in the Policy Name field. From the drop-down list, choose the switch on which this policy was deployed. The Deployment History pane displays the following fields. ▪ 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 the 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. ▪ 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 Day MMM DD YYYY HH:MM:SS Timezone. ▪ Failed Reason - Specifies why the policy was not successfully deployed.
---------------------------	--

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.
Host Role	Specifies the host device role. The host device role is one of the following: ▪ Sender ▪ Receiver ▪ Receiver-External ▪ Receiver-Local
Operation	Specifies if the operation of the host policy. The policy has the following operations: ▪ 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. ▪ 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** page.

Deployment Status Fields and Descriptions

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.

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

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Hosts > Host Policies**.
- Choose **Manage > 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:

1. In the **Host Policies** window, from the **Actions** drop-down list, choose **Create Host Policy**.
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**.



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.



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.



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.



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.
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

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Hosts > Host Alias**.
- Choose **Manage > Fabrics**. Double-click on a fabric to open **Fabric Overview > Hosts > Host Alias**.



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.

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 .

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 .csv file that contains the host IP address and corresponding unique host name information. Click Open. The host aliases are imported and displayed in the *Host Alias* window.
Export	Allows you to export host aliases for devices in the fabric. To export a host alias, choose Export. Select a location on your local system directory to store the host aliases configuration from Nexus Dashboard Fabric Controller and click Save. The host alias configuration file is exported to your local directory. The file name is appended with the date and time at which the file was exported. The format of the exported file is .csv.

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

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Hosts > Host Alias**.
- Choose **Manage > 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:

1. In the **Host Alias** window, from the **Actions** drop-down list, choose **Create Host Alias**.
2. In the *Create Host Alias* window, enter the following:



All the fields are mandatory.

- **VRF** - Select the VRF from this drop-down list. The default value is **default**.



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.



You can also create host alias before a host sends any data to its directly connected sender or receiver leaf.

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 **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Hosts > Applied Host Policies**.
- Choose **Manage > 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.

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: ▪ 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



This tab is only available on IPFM fabric when you have deployed IPFM on Nexus Dashboard Fabric Controller.

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric Summary** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Hosts**.
- Choose **Manage > 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 [Overview](#).

The **Flows** tab comprises the following horizontal tabs:

Flow Status

1. Navigate to the **Manage > Fabrics** page.
2. Double-click on a fabric to open the **Fabric Overview > Flows > Flow Status** page.

IPFM and Generic Multicast Flow Status

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.

IPFM Flow Status

In previous releases of NDFC, you could view the IPFM flow status only at the Layer 3 boundary, as a physical interface, or as a switch virtual interface (SVI). Layer 2 belonging to the SVI was not visible. Beginning with the NDFC 12.2.1 release, visibility in the Layer 2 segment is possible after the SVI. You can identify the receiver connected Layer 2 interface.

View the Layer 2 port information and the Layer 3 SVI in the **Receiver Interface** column, or by clicking on the **active link** under the **Flow Link State** column on the **Fabric Overview > Flows > Flow Status** page.

If you click on the **active** link, you can view the Layer 2 port in the topology diagram with an updated tooltip and a table displaying the Layer 2 physical port.

You can also view the Layer 2 receiver port along with the SVI details by navigating to the **Overview > Topology** page.

Generic Multicast Flow Status

In the generic multicast mode, the switch reports the receiver interface IP address instead of the receiver endpoint IP address. This IP is displayed in the **Flow Status** and **Topology** pages 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 display **allowed bytes/packets** and not **denied bytes/packets**.

Multicast Traffic Conversion

Field	Description
MUNAT	Specifies that the multicast traffic at the egress interface is converted to unicast traffic at the receiver interface.
Umcat	Specifies that the received multicast traffic at the egress interface is converted into unicast traffic at the sender interface.

1. Click the **Unicast** or **Multicast** link in the **Receiver** or the **Sender Interface** columns to view the IP route table for this interface.
2. Click the **active** link in the **Flow Link State** column to view the details for a given flow such as all pre or 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 for a particular multicast IP.

In VXLAN TRM, sources and receivers associated with an overlay flow are in a customer, also known as a tenant VRF. This tenant traffic is encapsulated in an underlay header that has **Encap Source** and **Encap Group** (located in the default VRF) on the sender VTEP side. The underlay encapsulated flow then reaches the receiver VTEP and is decapsulated here.

The flow topology in NDFC shows the overlay and underlay parts of the flow in different colors (purple for the underlay and green for the overlay).

Separation Between Default and Tenant VRFs

Field	Description
Type	Specifies the name of the virtual routing and forwarding (VRF).
L3VNI	Specifies the tenant VNI.
Encap Source	Specifies the IP address of the encapsulated source from the default VRF.
Encap Group	Specifies the IP address of the encapsulated group from the default VRF.

3. Click on the **Telemetry Sync Status** link above the table on the top-right corner.

The **Telemetry Sync Status** page displays the sync status and the IP address of the telemetry collector for each switch, along with the timestamp at the last sync.

4. To view the load on each telemetry collector, use the **Telemetry Collector == <IP address of the collector>** filter.

You can balance the collector performance based on the flows it is currently handling.

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 the same group. Nexus Dashboard Fabric Controller aggregates these flows per sender and receiver combinations and provides visibility into NAT rules through a topology. For more information about flow topology for active flows, see [RTP/EDI Flow Monitor](#).

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 or 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 or multicast group and filter relevant entries. Note that a searched IP address may not be visible in the main table on filtering as it may be part of a pre or a post entry that can be seen on the corresponding pop-up window.

For NAT flows with NAT types containing **Ingress**, the source and group will be the postNAT source and the postNAT group. For NAT types containing **Egress**, the source and group will be the preNAT source and the preNAT 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 the **NAT** label on the link to the receiver for the egress NAT.

For NAT flows, 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 a prefix. If the VRF is **default**, it will not be displayed.

NAT Fields and Descriptions

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 the First Hop Router (FHR). • Ingress (R) - Specifies that ingress NAT is performed on the Receiver Switch (also known as the Last Hop Router (LHR). • Ingress (S, R) - Specifies that ingress NAT is performed on both the Sender and the 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.
----------------------	---

Active Tab Fields and Descriptions

Field	Description
Common Fields for IPFM and Generic Multicast Modes	
VRF	Specifies the name of the VRF for the flow.
Encap	Specifies the name of the encap for the TRM flow.
Multicast IP	Specifies the multicast IP address for the flow.  You can click the wave link next to the multicast IP address to view the pictorial representation of the flow statistics.
Flow Alias	Specifies the name of the flow alias.
Flow Link State	Specifies the state of the flow link. 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 the receiver interfaces. 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 a spine.
Receiver Interface	Specifies the interface to which the receiver is connected to. Displays the Layer 2 physical port for the receiver interface. Example: Vlan120:Ethernet1/3
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.
Fields Specific for Generic Multicast Mode	
Receiver Interface IP	Specifies the IP address of the receiver interface joining the group.

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.  You can click the chart link next to the Multicast IP address to view the pictorial representation of the 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.
Policed	Specifies whether a flow is policed or not policed.
Receiver	Specifies the IP address or the host alias of the receiver joining the group.
Bandwidth	Specifies the bandwidth that is allotted for the traffic.
QOS/DSCP	Specifies the switch-defined QoS Policy.

Policy ID	Specifies the policy ID applied to the multicast IP.
Fault Reason	Specifies the reason for the inactive flow. Cisco Nexus Dashboard Fabric Controller determines the inactive flow if both the sender and receiver mroute exists with any of the following combinations. Options are: • Receiver IIF is null • Receiver OIF is null • Sender IIF is null • Sender OIF is null In this scenario, the switch will not have any fault reason. Therefore, there is no fault reason for such an inactive flows.
Field Specific for Generic Multicast Mode	
Receiver Interface IP	Specifies the IP address of the receiver interface joining the group.

Sender Only *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.
Flow Alias	Specifies the name of the flow alias.
Flow Link State	Specifies the flow link state, if it's allow or deny. Click the Sender Only link to view the network diagram or topology of the sender and the 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 the 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 name of the sender.
NAT	Specifies whether the flow is ingress, egress, or both ingress and egress.

Sender Switch	Specifies the IP address of the sender switch.
Sender Ingress Interface	Specifies the name of the sender ingress interface.
Sender Start Time	Displays the time from when the sender switch is transmitting information.
Fields Specific for IPFM Mode	
Policed	Specifies whether a flow is policed or not policed.
Policy ID	Specifies the policy ID applied to the multicast IP.
Bandwidth	Specifies the bandwidth that is allotted for the traffic.
QOS/DSCP	Specifies the switch-defined QoS Policy.
Priority	Specifies the flow priority for flows.

Receiver Only *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.
Flow Alias	Specifies the name of the flow alias.
Flow Link State	Specifies the flow link state, if it's allow or deny. Click the Receiver Only link to view the network diagram or topology of the sender and the 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 the 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.
Source Specific Sender	Specifies the IP address of the multicast sender.
Receiver	Specifies the receiver ID. If the multicast receiver is remote, the Remote label can be seen next to its name.
NAT	Specifies whether the flow is ingress, egress, or both ingress and egress.
Receiver Switch	Specifies the IP address of the receiver switch.

Receiver Interface	Specifies the name of the destination switch interface.
Receiver Join Time	Specifies the time at which the receiver joined.
Fields Specific for IPFM Mode	
Bandwidth	Specifies the bandwidth that is allotted for the traffic.
Policy ID	Specifies the policy ID applied to the multicast IP.
Priority	Specifies the flow priority for flows.
QOS/DSCP	Specifies the switch-defined QoS Policy.



If sstatistics 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 the statistical data in either a .csv or .pdf format.



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 high-availability 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 high-availability 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

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Hosts > Flow Policies**.
- Choose **Manage > Fabrics**. Double-click on a fabric to open **Fabric Overview > Hosts > Flow Policies**.

Use this window to configure the flow policies.



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.



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.



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.

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. • 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.  In adding or editing a flow policy, the default policer state is Enabled.
Last Updated	Specifies the date and time at which the flow policy was last updated. 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.



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.

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 .

Edit Flow Policy	Allows you to view or edit the selected flow policy parameters.  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.
Delete Flow Policy	Allows you to delete the user-defined flow policy.  * 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. 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.

Purge	Allows you to delete all the flow policies at a single instance.  Undeploy policies from all switches before deleting them from Nexus Dashboard Fabric Controller. To delete all flow policies, choose the Purge action. A warning message appears asking you to undeploy policies from all the switches. Click Confirm to proceed with deletion and leave the policies on the switches or click Cancel to discard the delete operation.
Import	Allows you to import flow policies from a csv file.  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. After import, all policies imported from a csv file are applied to all managed switches automatically. To import the flow policies, choose the Import action. Browse the directory and select the .csv file that contains the flow policy configuration information. The policy will not be imported if the format in the .csv file is incorrect. Click Open. The imported policies are automatically deployed to all the switches in the fabric.
Export	Allows you to export flow policies to a csv file. To export the flow policies, choose the Export action. Select a location on your local system directory to store the flow policy details file. Click Save. The flow policy file is exported to your local directory. The file name is appended with the date on which the file is exported. The format of the exported file is .csv.

Deploy Selected Policies	Select this option to deploy only the selected policies to the devices. You can deploy other policies when required. Select one or more check boxes next to the policy name. Select this option to deploy selected policies to the switch.
Deploy All Custom Policies	Select this option to deploy all the custom or user-defined policies at a single instance. 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: • 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. • Action - Specifies the action that is performed on the switch for that flow 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 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.

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



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:

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

The **Create Flow Policy** window is displayed.

2. In the **Create Flow Policy** window, specify the parameters in the following fields.
 - o **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**.



- 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.

- o **Policy Name** - Specify a unique policy name for the flow policy.
 - o **Bandwidth** - Specifies the bandwidth that is allocated for the flow policy. Select of the radio buttons to choose **Gbps**, **Mbps**, or **Kbps**.
3. From the **QoS/DSCP** drop-down list, choose an appropriate ENUM value.
 4. Click the **Policer** check box to enable or disable policer for a flow.
 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:

- o 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.

- o 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.

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

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Flows > Flow Alias**.
- Choose **Manage > Fabrics**. Double-click on a fabric to open **Fabric Overview > Flows > Flow Alias**.

Use this tab to configure flow alias.



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.

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.

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 .
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 .csv file that contains the flow IP address and corresponding unique flow name information. Click Open. The flow aliases are imported and displayed in the Flow Alias window.
Export	Allows you to export flow aliases for devices in the fabric. To export a flow alias, choose Export. Select a location on your local system directory to store the flow aliases configuration from Nexus Dashboard Fabric Controller and click Save. The flow alias configuration file is exported to your local directory. The file name is appended with the date and time at which the file was exported. The format of the exported file is .csv.

Creating Flow Alias

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Flows > Flow Alias**.
- Choose **Manage > 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:

1. In the **Flow Alias** window, from the **Actions** drop-down list, choose **Create Flow Alias**.
2. In the **Create Flow Alias** window, enter the following:



All the fields are mandatory.

- o **VRF** - Select the VRF from this drop-down list. The default value is **default**.



Host and IP Address are unique per VRF, that is, same host name with the same IP Address can exist in multiple VRFs.

- o **Flow Name** - Enter a fully qualified unique flow name for identification of the flow alias.
- o **Multicast IP Address** - Enter the multicast IP address for the flow alias.
- o **Description** - Enter a description for the flow alias.

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

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Hosts > Static Flow**.
- Choose **Manage > 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.

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 .
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.

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

Before you begin:

Select a switch in the **Static Flow** tab of the **Fabric Overview** window before creating a static flow for it.

To create a static flow for the selected switch, perform the following steps:

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

The **Create Static Flow** window is displayed.

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.

3. Click **Save & Deploy** to save the static flow.

Click **Cancel** to discard it.

Metrics

The **Metrics** tab displays the infrastructure health and status. You can view CPU utilization, Memory utilization, Traffic, Temperature, Interface, Links, and Power details.



Metrics tab is available only if you have enabled **Performance Monitoring** feature in Cisco Nexus Dashboard Fabric Controller.

Perform the following steps to navigate to the **Metrics** tab:

1. In Cisco Nexus Dashboard Fabric Controller, choose **Manage > Fabrics**.
2. Double-click the fabric name to open the **Fabric Overview** page.
3. Click on the **Metrics** tab.

The following table describes the columns that appear on the **CPU** and **Memory** tab.

Fields	Descriptions
Switch Name	Specifies the name of the switch.
IP Address	Specifies the switch IP address.
Low Value (%)	Specifies the lowest CPU/memory utilization value on the switch.
Avg. Value (%)	Specifies the average CPU/memory utilization value on the switch.
High Value (%)	Specifies the high CPU/memory 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 appear on the **Traffic** tab.



Inbound discards data is not available for Cisco Catalyst switches.

Fields	Descriptions
Switch Name	Specifies the name of the 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.

Fields	Descriptions
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.



Temperature specifications for Cisco Catalyst switches are currently not available.

Fields	Descriptions
Switch Name	Specifies the name of switch.
IP Address	Specifies the switch IP address.
Temperature Module	Specifies the module of temperature.
Low Value ©	Specifies the lowest temperature value.
Avg. Value ©	Specifies the average temperature value.
High Value ©	Specifies the high temperature value.
Show last day	Click Show last day to view data for selected day, week, month, and year.



Beginning with NDFC release 12.2.1, for the **Interface** and **Link** tabs described below, you can also filter the performance data that is shown using the following options:

- **Real time:** Gathers performance data every 10 seconds
- **Custom:** Gathers performance data based on the calendar begin and end dates that you select

The following table describes the columns that appears on **Interface** tab.

Fields	Descriptions
Switch	Specifies the name of the switch.
Interface	Specifies the name of the interface
Description	Specifies the description of the 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.

Fields	Descriptions
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 inbound average error value.
Out Avg.	Specifies the outbound peak error value.
In Peak	Specifies the inbound peak error value.
Out Peak	Specifies the outbound peak error value.
Discards	
In Avg.	Specifies the average discard value.
Out Avg.	Specifies the peak discard value.
In Peak	Specifies the inbound peak discard value.
Out Peak	Specifies the outbound 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 **Links** tab.

Fields	Descriptions
Switch	Specifies the name of the switch.
Speed	Specifies the switch speed.
Status	Specifies the status of switch.
Rx.	
Avg.	Specifies the average Rx value.
Avg%	Specifies the average percentage of Rx value.
Peak	Specifies the peak Rx value.
Peak%	Specifies the peak percentage Rx value.
Tx.	
Avg.	Specifies the average Tx value.
Avg%	Specifies the average percentage of Tx value.
Peak	Specifies the peak Tx value.

Fields	Descriptions
Peak%	Specifies the peak percentage Tx value.
Rx+Tx	Specifies the sum value of Rx and Tx.
Errors	
In Avg.	Specifies the inbound average error value.
Out Avg.	Specifies the outbound peak error value.
In Peak	Specifies the inbound peak error value.
Out Peak	Specifies the outbound peak error value.
Discards	
In Avg.	Specifies the average discard value.
Out Avg.	Specifies the peak discard value.
In Peak	Specifies the inbound peak discard value.
Out Peak	Specifies the outbound 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 appear on the **Power** tab.

Field	Description
Switch Name	Specifies the name of the switch. Click the name of a switch to access a Switch slide-in pane with the switch details.
IP Address	Specifies the switch IP address.
Power Module	Specifies the power module. Click the graph icon to view the power usage chart.
Power Usage (%)	Specifies the power usage as a percentage with subcategories based on the selected switch.
Avg	Specifies the average power usage.
Min	Specifies the minimum power usage.
Max	Specifies the maximum power usage.
Capacity (AMPs)	Specifies the capacity in (amps) amperes.
Draw (AMPs)	Specifies the power usage draw with subcategories based on the selected switch.
Avg	Specifies the average power usage draw.
Min	Specifies the minimum power usage draw.
Max	Specifies the maximum power usage based.
Capacity (WATTS)	Specifies the power capacity in watts.

Draw (WATTS)	Specifies the power usage draw in watts with subcategories based on the selected switch.
Avg	Specifies the average power usage draw in watts.
Min	Specifies the minimum power usage draw in watts.
Max	Specifies the maximum power usage draw in watts.

Click on **Day**, **Week**, or **Month** to filter the data based on time increments.

If you click on the **Day** increment, the graph displays one line for the selected data type. If you click on the **Week** or **Month** option, the graph displays the average, minimum, or maximum power usage values depending on the chosen data type. Click the drop-down list to view power data for the selected **Hour**, **Day**, **Week**, or **Month**.



The label for the drop-down list changes based on your selection. For example, if you clicked **Hour**, the label changes to **Show last hour**.

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.



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 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 **Manage > 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.
Undeploy All NAT Modes	Choose Undeploy All NAT Modes to undeploy all modes from the switch.
Redo All Failed NAT Modes	Choose Redo All Failed NAT Modes to deploy all failed modes.
Deployment History	Select a mode from the table and choose Deployment History to view the deployment history of the selected mode. Deployment History shows the following fields. ▪ 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

1. Choose **Manage > Fabrics**.
2. Double-click a fabric name.

The **Fabric Overview** window appears.

3. Click the **Multicast NAT** tab.
4. Click the **NAT Modes** tab.
5. Click **Actions > Create NAT Mode** to add a NAT mode.

The **Add NAT Mode** window appears.

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.

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.

7. Click **Save & Deploy** to save the NAT mode and deploy it.

Deleting a NAT Mode

1. Choose **Manage > Fabrics**.
2. Double-click a fabric name.

The **Fabric Overview** window appears.

3. Click the **Multicast NAT** tab.
4. Click the **NAT Modes** tab.
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.

6. Click **Confirm** to delete the selected NAT mode.

Recirc Mappings

NDFC allows you to map recirculation packets across ports for ingress or egress interfaces. From Release 12.1.1e, you can configure recirc mappings for the following translation types:

- Multicast-to-Multicast
- Multicast-to-Unicast
- Unicast-to-Multicast

Choose **Manage > Fabrics**. Double-click a fabric name and click **Multicast NAT > Recirc Mappings** to configure recirc mappings.

The following table describes the fields that appear on the **Recirc Mappings** tab.

Field	Description
VRF	Specifies the VRF over which the recirc mapping is routed.

Egress Interfaces	Specifies the egress interfaces for the mapping.
Destination/Prefix	Specifies the IP address of the destination unicast interface
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 **Recirc Mappings** tab.

Action Item	Description
Create NAT Recirc Mapping	Choose Create NAT Recirc Mapping to add an Recirc mapping.
Edit NAT Recirc Mapping	Select a mode from the table and choose Edit NAT Recirc Mapping to edit an Recirc mapping.
Delete NAT Recirc Mapping	Select a mode from the table and choose Delete NAT Recirc Mapping to delete an Recirc 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 Recirc mappings from Nexus Dashboard Fabric Controller to a CSV file.

Deploy Selected NAT Recirc Mappings	Select modes from the table and choose Deploy Selected NAT Recirc Mappings to deploy selected Recirc mapping to the switch.
Deploy All NAT Recirc Mappings	Choose Deploy All NAT Recirc Mappings to deploy all Recirc mappings to the switch.
Undeploy Selected NAT Recirc Mappings	Select modes from the table and choose Undeploy Selected NAT Recirc Mappings to undeploy selected Recirc mappings from the switch.
Undeploy All NAT Recirc Mappings	Choose Undeploy All NAT Recirc Mappings to undeploy all Recirc mapping from the switch.
Redo All Failed NAT Recirc Mappings	Choose Redo All Failed NAT Recirc Mappings to deploy all failed Recirc mappings.
Deployment History	Select a Recirc Mapping from the table and choose Deployment History to view the deployment history of the selected Recirc mapping. Deployment History shows the following fields. ▪ Switch Name—Specifies the name of the switch that the mode was deployed to. ▪ VRF—Specifies the VRF used to configure the selected recirc mapping. ▪ Map Interface—Specifies the map interface for the Recirc mappings. ▪ Max Replications—Specifies the maximum replications for the Recirc mappings. ▪ Egress Interfaces*or*Destination/Prefix—Specifies the interface over which Recirc mapping is configured. ▪ Deployment Status—Displays the status of deployment. It shows if the deployment was Success or Failed. If failed, the reason is displayed. ▪ Action—Specifies the action that is performed on the switch for that Recirc 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.

Adding Recirc Mapping

1. Choose **Manage > Fabrics**.
2. Double-click a fabric name.

The **Fabric Overview** window appears.

3. Click the **Multicast NAT > Recirc Mappings** tab.
4. From the **Selected Switch** drop-down list, select switch on which you want to create recirc mappings.
5. Click **Actions > Create Recirc Mapping** to add a recirculation mapping for the selected switch.

The **Add Recirc Mappings** window appears.

6. In the **Add Recirc Mappings** window, **Selected Switch** field specifies the switch name.

This field is read-only, and it's based on the switch selected in the Recirc Mappings window.

7. From the **VRF** drop-down list, select the vrf over which the recirc is routed.
8. In the Translation Type, select one of the translation types:
 - Multicast-to-Multicast
 - Multicast-to-Unicast
 - Unicast-to-Multicast
9. If you selected **Multicast-to-Multicast** transition type, in the **Egress Interfaces** area, select one of the following:
 - All - Choose All to select all the interfaces
 - Select one or more - 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.
10. Based on the transition type, do the following:
 - If you selected **Multicast-to-Unicast** transition type, enter the IP address of the destination unicast interface in the **Destination/Prefix** field.
 - If you selected **Unicast-to-Multicast** transition type, enter the IP address of the destination multicast interface in the **Destination/Prefix** field.
11. From the **Map Interface** drop-down list, select an interface to start recirc mapping.

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.

12. In the **Max Replications** field, enter the maximum replications for the map interface. The range for this field is 1-40. The default value is 40.
13. Click **Save & Deploy** to save the NAT mode and deploy it.

NAT Rules

NAT Rules NAT rules are identical for ingress and egress NAT except you need to also specify receiver OIF for egress NAT.

Choose **Manage > 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.

Deployment History	Select a rule from the table and choose Deployment History to view the deployment history of the selected rule. Deployment History shows the following fields. ▪ 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

1. Choose **Manage > Fabrics**.
2. Double-click a fabric name.

The **Fabric Overview** window appears.

3. Click the **Multicast NAT** tab.
4. Click the **NAT Rules** tab.
5. Click **Actions > Create NAT Rule** to add a NAT rule.

The **Add NAT Rule** window appears.

6. In the **Add NAT Rule** window, specify the following information:

Translation Type: Select one of the translation types:

- Multicast-to-Multicast
- Multicast-to-Unicast
- Unicast-to-Multicast

Mode: Select the NAT mode, that is, **Ingress** or **Egress**.

This mode is not visible for Multicast-to-Unicast and Unicast-to-Multicast translation types.

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/Unicast IP: Specifies the multicast or unicast group before NAT.

Post-Translation Group: Specifies the multicast or unicast 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.



The Post-Translation Source IP needs to be the secondary IP address on the loopback interface to make sure RPF check won't fail. However, the switch maintains separate records for Pre- and Post- NAT records, and NDFC merges unicast-multicast pre-post entries as single flow.

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.

Static Oif: This field is not visible for Ingress mode. In Egress mode, this field displays **Egress Interfaces** defined in the Recirc Mappings screen. The field is empty if there are no mappings defined.

7. Click **Save & Deploy** to save the NAT rule and deploy it.

Deleting NAT Rule

1. Choose **Manage > Fabrics**.
2. Double-click a fabric name.

The **Fabric Overview** window appears.

3. Click the **Multicast NAT** tab.
4. Click the **NAT Rules** tab.
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.

6. Click **Confirm** to delete the selected NAT rule.

RTP/EDI Flow Monitor



This tab is only available on an IPFM fabric when you have deployed IPFM on the Nexus Dashboard Fabric Controller.

1. Click **Manage > Fabrics**.
2. Click on a fabric to open the **Fabric** slide-in pane.
3. Click the **Launch** icon.
4. Click **Fabric Overview > RTP/EDI Flow Monitor**.
5. Click **Manage > Fabrics**.
6. Double-click on a fabric to open **Fabric Overview > RTP/EDI Flow Monitor**.



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.



You need to enable telemetry on the switches to view the RTP/EDI Flow Monitor. For more information, see your respective platform documentation.

The description of the fields in these tabs are:

Field	Description
VRF	Specifies the name of the VRF.
Switch	Specifies the name of the switch.
Interface	Specifies the interface from which the flows are detected.
Source IP	Specifies the source IP address of the flow.
Source Port	Specifies the source port of the flow.
Destination IP	Specifies the destination IP address of the flow.
Destination Port	Specifies the destination port of the flow.
Receiver IPs	Specifies the receiver IPs which are connected directly to the given switch.
Bit Rate	Specifies the bit rate of the flow, in bps, kbps, mbps, gbps, or tbp.
Packet Count	Specifies the number of packets in the flow.
Packet Loss	Specifies the number of lost packets.
Loss Start	Specifies the time at which the packet loss started.

Loss End	Specifies the time at which the packet loss stopped.
Start Time	Specifies the time at which the flow started.
Protocol	Specifies the protocol that is being used for the flow.

You can click the **Telemetry Switch Sync Status** link to check whether the switches are in sync. The **Telemetry Sync Status** page 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** page 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 on the **Flow Status** page. For more information about multicast NAT visualization, see [Flow Status](#).

From Cisco NDFC Release 12.1.2e, the flow topology for the active flows is displayed on the **Active Flows** tab.

1. 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 **®** 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.

2. Hover your cursor over a switch to display the following details:
 - Name
 - IP address
 - Model
 - Packet loss, if any
3. 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-up dialog.

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 seven days. You can customize this setting by entering the required value in the **IPFM history retention days** field in **Fabric Controller > Admin > System Settings > Server Settings > IPFM** and saving it.



The **Drop History** tab displays only the last 100,000 records at the maximum.

Global Config



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).

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Global Config**.
- Choose **Manage > 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

- Choose **Manage > 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 **Manage > 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.



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.

Switch Global Config Table Fields and Description

Field	Description
VRF	Specifies the name of the VRF. This VRF is used to associate IPFM Host/Flow policies as well as Host/Flow aliases for both IPFM and Generic Multicast fabrics.
Unicast Bandwidth Reservation %	Displays a numeric value that indicates the unicast bandwidth configuration percentage, and the status specifies if the bandwidth deployment was success, or failed or not deployed. You can configure the server to allot a dedicated percentage of bandwidth to unicast traffic. The remaining percentage is automatically reserved for multicast traffic. 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. 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.

Reserve Bandwidth to Receiver Only	Bandwidth reservation status specifies if the bandwidth deployment was success, or failed or not deployed. 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. 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. 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.
ASM/MASK	Displays the number of Any Source Multicast (ASM) groups enabled for the selected VRF and the status indicates whether the ASM and Mask configuration was deployed successfully, or failed or not deployed. 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. The IP address and subnet mask in the ASM/MASK field define the multicast source. The ASM range is configured by specifying the IP address and the subnet mask. 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. 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.

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.

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.

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.

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

- Choose **Manage > Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview > Global Config > IPFM VRF**.
- Choose **Manage > Fabrics**. Double-click on a fabric to open **Fabric Overview > Global Config > IPFM VRF**.

From Cisco NDFC Release 12.1.2e, **IPFM VRF** tab is included under the **Fabric Overview** window. Choose **LAN > Fabrics**. Double-click on a fabric to open **Fabric Overview > IPFM VRF**.

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

From Cisco NDFC Release 12.1.2e, you can configure and monitor both NBM active and passive VRFs. In NBM passive mode, NDFC will be involved only in the monitoring of IPFM fabric and not configuration except in setting up VRF mode as NBM passive. Perform the following steps to change the NBM mode:

- Click **Actions > Create VRF**.
- On the **Create VRF** window, enter the name of the VRF. Choose **Active** or **Passive** and click **Save & Deploy**.



You cannot edit the existing VRF to change the NBM mode. You must delete and re-create VRF to change the NBM mode from active to passive or conversely. If fabric is set to monitor mode, changing VRF is not applicable as this is fabric level configuration and not VRF configuration.

You are not allowed to create **IPFM VRF** when none of the switches are imported to NDFC. Import or add switch to the fabric to create IPFM VRF.

Discovery status is updated at regular interval by a background process. NBM configuration can be deployed even if the switch is in an unreachable state. After periodic discovery, the status of switches are updated appropriately.

IPFM VRF Table Fields and Description

Field	Description
Name	Specifies the name of the VRF.
Mode	Specifies the type of mode (Active or Passive) of the VRF.
Deployment Status	Specifies whether the VRF deployment is successful, failed, or the VRF is not deployed. For default VRFs, the deployment status is displayed as Not Applicable . Click the Failed status to view more information about the Switch Global Config .
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.

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, choose Active or Passive mode and click Save & Deploy to retain the changes and deploy or click Cancel to discard the changes.  When you create an active nondefault VRF, although the default host and flow policies are automatically created for that VRF, you must manually deploy the policies to the switches in the fabric. When VRF is set to passive, then flow policies are not created. For more information about deploying the policies manually, see Host Policies 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.

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.

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)



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.

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

Use the **VRF** window to create, edit, and delete VRFs.

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.

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.
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.

Virtual Infrastructure

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

Field	Description
VM Name	Specifies the name of the virtual machine.
IP Address	Specifies the IP address of the virtual machine.
MAC Address	Specifies the MAC address of the virtual machine.
VLAN	Specifies the VLAN associated with the virtual machine.
Network	Specifies the network associated with the virtual machine.
VRF	Specifies the VRF associated with the virtual machine.
Security Group	Introduced in NDFC release 12.2.2. Specifies the security group associated with the virtual machine. For more information, see Configuring Security for VXLAN EVPN Fabrics .
Switch	Specifies the switch connected to the virtual machine.
Switch Interface	Specifies the switch interface connected to virtual machine.
State	Specifies the state of virtual machine.
Alerts	Displays any alerts associated with the virtual machine.

Beginning with NDFC release 12.2.2, you can also associate a VM UUID or VM NIC port with a security group. For more information, see [Configuring Security for VXLAN EVPN Fabrics](#).

1. In the Virtual Infrastructure window, select the VM that you want to associate with a security group.
2. Click **Actions > Set Group ID**.

The **Set Security Group ID** window appears.

3. Select an existing security group, or click **Add Security Group**.

See "Create a Security Group" in [Configuring Security for VXLAN EVPN Fabrics](#) for more info.

4. Click **Submit**.

This action configures an IP selector for the VRF in the switches.

5. Complete any other necessary configurations.
 - o To remove an association with a security group on a VM, select that VM and click **Actions > Remove Group ID**.
 - o When you have completed all necessary configurations related to security groups, including setting or removing group IDs in this page, perform a **Recalculate & Deploy** to generate the corresponding intent.

Copyright

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

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

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

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

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

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

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

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/go/trademarks>. 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. (1110R)

© 2017–2025 Cisco Systems, Inc. All rights reserved.