



# Editing Campus VXLAN Fabric Settings, Release 4.2.1

# Table of Contents

New and changed information	1
About Campus VXLAN EVPN fabric	2
Editing Campus VXLAN fabric settings	3
General	4
Fabric Management	5
General Parameters	5
Replication	7
Protocols	7
Advanced	7
Resources	9
Bootstrap	12
Configuration Backup	14
Border Gateway	14
Telemetry	17
Edit configuration settings	18
Edit NAS settings	18
Disable Telemetry	19
Perform force disable telemetry on your fabric	19
NAS	19
Guidelines and limitations for network attached storage	20
Add network attached storage to export flow records	20
Add NAS to Nexus Dashboard	20
Add the onboarded NAS to Nexus Dashboard	21
Flow collection	23
Understanding flow telemetry	23
Flow telemetry guidelines and limitations	23
Configure flows	25
Monitor the subnet for flow telemetry	27
Understanding Netflow	29
Understanding Netflow types	29
Netflow guidelines and limitations	30
Configure Netflow	31
Understanding sFlow	32
Guidelines and limitations for sFlow	32
Configure sFlow telemetry	32
External streaming	34
Configure external streaming settings	34
Guidelines and limitations	35
Guidelines and limitations in NX-OS fabrics	35
Email	35
Message bus	37

Add Kafka broker configuration . . . . .	37
Configure Kafka exports in fabric settings . . . . .	39
Anomalies . . . . .	41
Advisories . . . . .	42
Statistics . . . . .	42
Faults . . . . .	43
Audit Logs . . . . .	43
Syslog . . . . .	43
Guidelines and limitations for syslog . . . . .	44
Add syslog server configuration . . . . .	44
Configure syslog to enable exporting anomalies data to a syslog server . . . . .	44
Additional settings . . . . .	47
Layer 3 VNI without VLAN . . . . .	47
Guidelines and limitations for Layer 3 VNI without VLAN . . . . .	47
Configuring automatic generation of BGP neighbor description . . . . .	48
Guidelines and limitations for automatic generation of a BGP neighbor description . . . . .	48
How to configure automatic generation of a BGP neighbor description . . . . .	48
How to configure automatic generation of a BGP neighbor description between border gateways in a VXLAN fabric group . . . . .	49
Adding Cisco Catalyst 9000 series switches and Nexus 9000 series switches to a Campus VXLAN EVPN fabric . . . . .	49
Recalculating and deploying configurations . . . . .	52
Creating VRFs in Campus VXLAN EVPN Fabric . . . . .	52
Attaching VRFs to switches in Campus VXLAN EVPN fabrics . . . . .	53
Creating and deploying networks in Campus VXLAN EVPN fabrics . . . . .	54
Creating networks for Campus VXLAN EVPN fabrics . . . . .	54
Attaching networks in a Campus VXLAN EVPN fabric . . . . .	57
Deploying networks in Campus VXLAN EVPN fabrics . . . . .	57
Creating DCI links for switches in Campus VXLAN EVPN fabrics . . . . .	57
Copyright . . . . .	59

# 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
Nexus Dashboard 4.2.1		There were no major changes from the previous release.

# About Campus VXLAN EVPN fabric

Nexus Dashboard supports Campus VXLAN EVPN fabric type to automate and manage Enterprise Campus VXLAN BGP EVPN networks based on Catalyst 9000 series switches. Optionally, administrators can integrate a Nexus 9000 switch with Border Gateway functionality to interconnect with remote Data Centers and Campus for VXLAN EVPN multi-fabric Layer 2 and Layer 3 extensions.

This document describes how to create a Campus VXLAN EVPN fabric with Cisco Catalyst 9000 series switches and Nexus 9000 series switches using the **Campus VXLAN EVPN** fabric template. This fabric supports OSPF as the underlay protocol and BGP EVPN as the overlay protocol. Using this fabric template, Nexus Dashboard manages all the configurations of a VXLAN EVPN fabric consisting of Cisco Catalyst 9000 IOS XE and Nexus 9000 NX-OS switches. Backing up and restoring this fabric is similar to **Data Center VXLAN EVPN** backup and restore.

Nexus Dashboard provides support for the following features with Campus VXLAN EVPN fabrics:

- Tenant Routed Multicast (TRM)
- Zero-Touch Provisioning (ZTP) on Cisco Catalyst switches through Plug and Play (PnP)
- Deploying Campus VXLAN EVPN fabric as a child fabric in VXLAN EVPN Multi-Site.

## Guidelines for configuring Campus VXLAN EVPN fabrics

- Out-of-band management is required for Campus VXLAN EVPN fabric types.
- Provides support for EVPN VXLAN Distributed Anycast Gateway when each SVI is configured with the same Anycast Gateway MAC.
- Provides support for Cisco Catalyst switches with Stackwise or Stackwise Virtual.
- Provides support for spine, leaf, and border roles on Cisco Catalyst switches. Whereas, Cisco Nexus 9000 series switches support border gateway, border gateway spine and border gateway super spine roles.
- Does not support Brownfield deployments.
- Does not support IPv6 underlay and Anycast RP.
- Does not support ISIS, ingress replication, unnumbered intra-fabric link, and 4 bytes BGP ASN.
- Does not support breakout interfaces on Cisco IOS XE switches.



For information about configuration compliance, see the section "Configuration Compliance in External Fabrics" in [Editing External Fabric Settings](#).

# Editing Campus VXLAN fabric settings

A **Campus VXLAN** fabric is a type of fabric that is used for a VXLAN EVPN campus deployment with Catalyst 9000 and Nexus 9000 switches as border gateways.

When you first create a Campus VXLAN fabric using the procedures provided in [Creating Fabrics and Fabric Groups](#), the standard workflow allows you to create a fabric using the bare minimum settings so that you are able to create a fabric quickly and easily. Use the procedures in this article to make more detailed configurations for your campus VXLAN fabric.

Perform the following steps to create the Campus VXLAN EVPN fabric for Cisco Catalyst 9000 series switches and Nexus 9000 series switches:

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

**Manage > Fabrics**

2. Locate the Campus VXLAN fabric that you want to edit.

Campus VXLAN fabrics are shown with **Campus VXLAN EVPN** in the **Type** column.

3. Click the circle next to the Campus VXLAN fabric that you want to edit to select that fabric, then click **Actions > Edit Fabric Settings**.

The **Edit *fabric\_name* Settings** page appears.

4. Click the appropriate tab to edit these settings for the fabric:
  - o [General](#)
  - o [Fabric Management](#)
  - o [Telemetry](#) (if the **Telemetry** feature is enabled for the fabric)

# General

Use the information in this section to edit the settings on the **General** page for your campus VXLAN fabric.

The fields on this **Edit Fabric Settings > General** page are the same fields that appeared on the **3. Settings** page when you first created this Campus VXLAN fabric, with the same configurations that you provided when you created this fabric. Edit these already-configured settings if necessary, or click another tab to leave these settings unchanged.

Change the general parameters that you configured previously for the Campus VXLAN fabric, if necessary.

<b>Fabric type</b>	<b>Description</b>
Name	The name for the fabric. This field is not editable.
Type	The fabric type for this fabric. This field is not editable.
Location	Choose the location for the fabric.
BGP ASN for Spines	Enter the BGP autonomous system number (ASN) for the fabric's spine switches.
License tier	Choose the licensing tier for the fabric: <ul style="list-style-type: none"><li>• <b>Essentials</b></li><li>• <b>Advantage</b></li><li>• <b>Premier</b></li></ul> Click on the information icon (i) next to License tier to see what functionality is enabled for each license tier.
Security domain	Choose the security domain for the fabric.

# Fabric Management

Use the information in this section to edit the settings on the **Fabric Management** page for your Campus VXLAN fabric. The tabs and their fields on the page are explained in these sections. The fabric-level parameters are included in these tabs.

- [General Parameters](#)
- [Replication](#)
- [Protocols](#)
- [Advanced](#)
- [Resources](#)
- [Bootstrap](#)
- [Configuration Backup](#)
- [Border Gateway](#)

## General Parameters

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

Field	Description
<b>Underlay Subnet IP Mask</b>	Specifies the subnet mask for the fabric interface IP addresses.
<b>Link-State Routing Protocol</b>	Specifies the supported routing protocol which is OSPF.

Field	Description
<b>Route-Reflectors</b>	<p>The number of spine switches that are used as route reflectors for transporting BGP traffic. Choose option 2 or 4 from the drop-down list. The default value is 2.</p> <p>To deploy spine devices as RRs, Nexus Dashboard sorts the spine devices based on their serial numbers and designates two or four spine devices as RRs. If you add more spine devices, existing RR configuration won't change.</p> <p>Increasing the count - You can increase the route reflectors from two to four at any point in time. Configurations are automatically generated on the other two spine devices designated as RRs.</p> <p>Decreasing the count - When you reduce four route reflectors to two, remove the required route reflector devices from the fabric.</p> <p>Follow these steps to reduce the count from 4 to 2.</p> <ol style="list-style-type: none"> <li>1. Change the value in the drop-down list to 2.</li> <li>2. Identify the spine switches designated as route reflectors. An instance of the <b>rr_state</b> policy is applied on the spine switch if it's a route reflector.</li> <li>3. Delete the spine devices that are not required from the fabric.</li> </ol> <p>If you delete existing RR devices, the next available spine switch is selected as the replacement RR.</p> <ol style="list-style-type: none"> <li>4. Click <b>Deploy Config</b> in the fabric topology window.</li> </ol> <p>You can preselect RRs and RPs before performing the first <b>Save &amp; Deploy</b> operation.</p>
<b>Anycast Gateway MAC</b>	Specifies the shared MAC address for the leaf switches.
<b>Enable Performance Monitoring</b>	<p>Enables performance monitoring on the switches.</p> <p>Ensure that you do not clear interface counters from the Command Line Interface of the switches. Clearing interface counters can cause the Performance Monitor to display incorrect data for traffic utilization. If you must clear the counters and the switch has both <b>clear counters</b> and <b>clear counters snmp</b> commands (not all switches have the <b>clear counters snmp</b> command), ensure that you run both the main and the SNMP commands simultaneously. For example, you must run the <b>clear counters interface ethernet slot/port</b> command followed by the <b>clear counters interface ethernet slot/port snmp</b> command. This can lead to a one time spike.</p>

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

# Replication

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

Field	Description
<b>Replication Mode</b>	Specifies the mode of replication used in the fabric for BUM (Broadcast, Unknown Unicast, Multicast) traffic. <b>Multicast</b> is selected by default.
<b>Multicast Subnet Group</b>	Specifies the IP address prefix used for multicast communication. A unique IP address is allocated from this group for each overlay network.  The replication mode change is not allowed if a policy template instance is created for the current mode. For example, if a multicast related policy is created and deployed, you cannot change the mode to Ingress.
<b>Enable Tenant Routed Multicast (TRM)</b>	Enables Tenant Routed Multicast (TRM) that allows overlay multicast traffic support over EVPN/MVPN in the VXLAN EVPN fabric.
<b>Rendezvous-Points</b>	Specifies the number of spine switches acting as rendezvous points.
<b>Underlay RP Loopback Id</b>	Specifies the loopback ID used for the RP, for multicast protocol peering purposes in the fabric underlay. The default is 254.

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

# Protocols

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

Field	Description
<b>Underlay Routing Loopback Id</b>	Specifies the loopback interface ID. By default, value 0 is populated as loopback0 that is normally used for fabric underlay IGP peering purposes.
<b>Underlay VTEP Loopback Id</b>	Specifies the loopback interface ID. By default, value 1 is populated as loopback1 used as the VTEP address.
<b>OSPF Process Id</b>	Specifies the OSPF process tag.
<b>OSPF Area Id</b>	Specifies the OSPF unique 32-bit area ID denoted in dotted decimal format.

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

# Advanced

The fields in the **Advanced** tab are described in the following table. All the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields

if needed.

Field	Description
<b>VRF Template</b>	Specifies the VRF template for creating VRFs. By default, the system uses the pre-defined <b>Default_VRF_Universal</b> template for overlay configuration for leaf switches.
<b>Network Template</b>	Specifies the network template for creating networks. By default, the system uses the pre-defined <b>Default_Network_Universal</b> template for leaf switches.
<b>VRF Extension Template</b>	Specifies the VRF extension template for enabling VRF extension to other fabrics. By default, the system uses the pre-defined <b>Default_VRF_Extension_Universal</b> template for border switches.
<b>Network Extension Template</b>	Specifies the network extension template for extending the network to other fabrics. By default, the system uses the pre-defined <b>Default_Network_Extension_Universal</b> template for border switches.
<b>Intra Fabric Interface MTU</b>	<p>Specifies the MTU for the intra fabric interface. The value must be an even number.</p> <p>The valid values range from 576 to 9216. This is a mandatory field.</p>
<b>Layer 2 Host Interface MTU</b>	<p>Specifies the MTU for the layer 2 host interface. This value must be an even number. The valid values range from 1500 to 9216.</p> <p>The default MTU for both fabric interface and host interface is 9198. The default MTU for IOS XE is 1500.</p>
<b>IOS XE System MTU</b>	Specifies the MTU for an IOS XE device. This value must be an even number. The valid values range from 1500 to 9198.
<b>Enable Tenant DHCP</b>	<p>Check the check box to enable DHCP and associated configurations globally on all the switches in the fabric. This is a prerequisite for supporting DHCP for overlay networks that are part of the tenant VRFs.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>Ensure that <b>Enable Tenant DHCP</b> is enabled before enabling DHCP-related parameters in the overlay profiles.</p> </div>
<b>VTP Mode</b>	By default, the VTP mode is <b>Off</b> . <b>Transparent</b> mode allows you to relay all VTP protocol packets that it receives on a trunk port to all other trunk ports.
<b>Enable NDFC as Trap Host</b>	Allows you to configure Nexus Dashboard as an SNMP trap destination.

Field	Description
<b>Enable Overlay Template Conversion</b>	<p>Allows you to convert all the existing VRFs and networks to use the default templates. In existing deployments using the <b>IOS_XE_VRF</b> and <b>IOS_XE_Network</b> templates, enabling this field converts the templates to use the <b>Default_VRF_Universal</b>, <b>Default_Network_Universal</b>, <b>Default_VRF_Extension_Universal</b> and <b>Default_Network_Extension_Universal</b> templates after performing a <b>Recalculate and deploy</b>.</p> <p>When adding child fabrics to a VXLAN fabric group, a Campus VXLAN EVPN fabric does not allow you to add fabrics that are configured with <b>IOS_XE_VRF</b> and <b>IOS_XE_Network</b> as the templates for existing VRFs and networks. Enable the <b>Enable Overlay Template Conversion</b> field to convert the existing VRFs and networks using <b>IOS_XE_VRF</b> and <b>IOS_XE_Network</b> templates to the <b>Default_VRF</b> and <b>Default_Network</b> templates.</p>
<b>Leaf Freeform Config</b>	Configures additional CLIs for all the Cisco Catalyst leaf switches in the fabric.
<b>Spine Freeform Config</b>	Configures additional CLIs for all the Cisco Catalyst spine switches in the fabric.
<b>Intra-fabric Links Additional Config</b>	Configures additional CLIs for all the intra fabric links.

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

## Resources

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

Field	Description
<b>Underlay Routing Loopback IP Range</b>	Specifies the loopback IPv4 addresses for protocol peering.
<b>Underlay VTEP Loopback IP Range</b>	Specifies the loopback IP address range for VTEPs.
<b>Underlay RP Loopback IP Range</b>	Specifies anycast or phantom RP IP address range.
<b>Underlay Subnet IP Range</b>	Specifies the IP addresses for underlay P2P routing traffic between interfaces.
<b>Layer 2 VXLAN VNI Range</b>	Specify the VXLAN VNI IDs for the fabric.
<b>Layer 3 VXLAN VNI Range</b>	Specify the VXLAN VNI IDs for the fabric.

Field	Description
<b>Network VLAN Range</b>	VLAN range for the per switch overlay network (min:2, max:4094).
<b>VRF VLAN Range</b>	VLAN range for the per switch overlay Layer 3 VRF (min:2, max:4094).
<b>Subinterface Range</b> <b>Dot1q Range</b>	Specifies the subinterface range when L3 sub-interfaces are used.
<b>VRF Lite Deployment</b>	<p>Specifies the VRF Lite method for extending inter-fabric connections.</p> <p>The <b>VRF Lite Subnet IP Range</b> field specifies resources reserved for IP address used for VRF Lite when VRF Lite IFCs are auto-created. If you select <b>Back2Back&amp;ToExternal</b>, then VRF Lite IFCs are auto-created.</p>
<b>Auto Deploy for Peer</b>	<p>This check box is applicable for VRF Lite deployment. When you select this checkbox, auto-created VRF Lite IFCs will have the <b>Auto Generate Configuration for Peer</b> field in the <b>VRF Lite</b> tab set, if the peer is a Cisco device.</p> <p>To access VRF Lite IFC configuration, navigate to the <b>Links</b> tab, select the link, and then choose <b>Actions &gt; Edit</b>.</p> <p>You can check or uncheck the check box when the <b>VRF Lite Deployment</b> field is not set to <b>Manual</b>. This configuration only affects the new auto-created IFCs and does not affect the existing IFCs. You can edit an auto-created IFC and check or uncheck the <b>Auto Generate Configuration for Peer</b> field. This setting takes priority always.</p>
<b>Auto Deploy Default VRF</b>	<p>When you select this check box, the <b>Auto Generate Configuration on default VRF</b> field is automatically enabled for auto-created VRF Lite IFCs. You can check or uncheck this check box when the VRF Lite Deployment field is not set to <b>Manual</b>. The <b>Auto Generate Configuration on default VRF</b> field when set, automatically configures the physical interface for the border device, and establishes an eBGP connection between the border device and the edge device or another border device in a different VXLAN EVPN fabric.</p>
<b>Auto Deploy Default VRF for Peer</b>	<p>When you select this check box, the <b>Auto Generate Configuration for NX-OS Peer on default VRF</b> field is automatically enabled for auto-created VRF Lite IFCs. You can check or uncheck this check box when the VRF Lite Deployment field is not set to <b>Manual</b>. The <b>Auto Generate Configuration for NX-OS Peer on default VRF field</b> when set, automatically configures the physical interface and the eBGP commands for the peer NX-OS and IOS XE switches.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 20px;">  <p>To access the <b>Auto Generate Configuration on default VRF</b> and <b>Auto Generate Configuration for NX-OS Peer on default VRF</b> fields for an IFC link, navigate to the <b>Links</b> tab, select the link and choose <b>Actions &gt; Edit</b>.</p> </div>
<b>Redistribute BGP Route-map Name</b>	Route Map used to redistribute BGP routes to IGP in default VRF for auto-created VRF Lite IFC links.

Field	Description
<b>VRF Lite Subnet IP Range</b>	<p>These fields are prefilled with the DCI subnet details. Update the fields as needed.</p> <p>The values shown on the page are automatically generated. If you want to update the IP address ranges, VXLAN Layer 2/Layer 3 network ID ranges or the VRF/network VLAN ranges, ensure that each fabric has its own unique range and is distinct from any underlay range to avoid possible duplication. You should only update one range of values at a time.</p>
<b>VRF Lite Subnet Mask</b>	<p>If you want to update more than one range of values, do it in separate instances. For example, if you want to update Layer 2 and Layer 3 ranges, you should do the following.</p> <ol style="list-style-type: none"> <li>1. Update the Layer 2 range and click <b>Save</b>.</li> <li>2. Click the <b>Edit Fabric</b> option again, update the Layer 3 range, and click <b>Save</b>.</li> </ol>
<b>Auto Allocation of Unique IP on VRF Extension over VRF Lite IFC</b>	<p>When enabled, IP prefix allocated to the VRF Lite IFC is not reused on VRF extension over VRF Lite IFC. Instead, unique IP Subnet is allocated for each VRF extension over VRF Lite IFC.</p>
<b>Per VRF Per VTEP Loopback Auto-Provisioning</b>	<p>Enables you to auto provision a loopback on a VTEP on VRF attachment.</p> <ol style="list-style-type: none"> <li>1. Enable the <b>Per VTEP Loopback Auto-Provisioning</b> option.</li> <li>2. Save the fabric settings.</li> <li>3. Perform a <b>Recalculate and deploy</b> operation.</li> <li>4. Navigate to <b>VRF Attachments</b>.</li> <li>5. If certain VRFs are already attached, click <b>Actions &gt; Quick attach</b>. This generates the new loopback in the VRF.</li> </ol> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> If VRF extensions are already enabled and configured, for example, VRF Lite on a border device, prior to enabling the fabric setting, you need to access the respective VRF attachment and the border device to reattach the VRF extension again. For example, VRF Corp is attached on Border-1 and extended to an external domain using VRF Lite. In this situation, when you perform a <b>Quick attach</b> to provision the new loopback in the VRF, the original VRF-Lite extension gets detached. You can then select the VRF attachment, edit, and re-attach the VRF-Lite extension and then deploy all the relevant configurations.</p> </div>
<b>Per VRF Per VTEP IP Pool for Loopbacks</b>	<p>Indicates the prefix pool to assign IP addresses to loopbacks on VTEPs on a per VRF basis.</p>

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

# Bootstrap

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

Field	Description
<b>Enable Bootstrap</b>	<p>Check this check box to enable the bootstrap feature. Bootstrap allows easy day-0 import and bring-up of new devices into an existing fabric. Bootstrap leverages POAP for NX-OS and PnP for IOS XE.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  For NX-OS switches, POAP will not work if you have set <b>Bootstrap Script Download Protocol</b> (in the <b>Server Settings</b> for LAN) as <b>https</b>.         </div>
<b>Enable Local DHCP Server</b>	<p>Select this check box to initiate enabling of automatic IP address assignment through a local DHCP server. When you select this check box, the <b>DHCP Scope Start Address</b> and <b>DHCP Scope End Address</b> fields become editable.</p> <p>If you want to configure a remote or external DHCP server for automatic IP address assignment, enter details about the external DHCP server in the <b>Switch Mgmt Default Gateway</b> and <b>Switch Mgmt IP Subnet Prefix</b> fields.</p>
<b>DHCP Version</b>	<p>Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the <b>Switch Mgmt IPv6 Subnet Prefix</b> field is disabled. If you select DHCPv6, the <b>Switch Mgmt IP Subnet Prefix</b> is disabled.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  Cisco Nexus 9000 and 3000 Series Switches support IPv6 POAP only when switches are either Layer-2 adjacent (eth1 or out-of-band subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other than /64 are not supported.         </div>
<b>Domain name</b>	Specifies the domain name of the DHCP server.
<b>DHCP Scope Start Address</b> and <b>DHCP Scope End Address</b>	Specifies the first and the last IP addresses of the IP address range to be used for the switch out of band POAP.
<b>Switch Mgmt Default Gateway</b>	Specifies the default gateway for the management VRF on the switch.

Field	Description
<b>Switch Mgmt IP Subnet Prefix</b>	<p>Specifies the prefix for the management interface on the switch. The prefix should be between 8 and 30.</p> <p><i>DHCP scope and management default gateway IP address specification:</i> If you specify the management default gateway IP address 10.0.1.1 and subnet mask 24, ensure that the DHCP scope is within the specified subnet, between 10.0.1.2 and 10.0.1.254.</p>
<b>Switch Mgmt IPv6 Subnet Prefix</b>	<p>Specifies the IPv6 prefix for the Mgmt0 interface on the switch. The prefix should be between 64 and 126. This field is editable if you enable IPv6 for DHCP.</p>
<b>Bootstrap Freeform Config (IOS-XE)</b>	<p>(Optional) Enter additional commands for IOS XE switches, as needed. For example, if you require some additional configurations to be pushed to the device and be available post device bootstrap, they can be captured in this field, to save the desired intent. After the devices boot up, they will contain the configuration defined in the <b>Bootstrap Freeform Config</b> field.</p> <p>Copy-paste the running-config to a <b>freeform config</b> field with correct indentation, as seen in the running configuration on the IOS XE switches. The freeform config must match the running config. For more information, see <a href="#">Enabling Freeform Configurations on Fabric Switches</a>.</p>
<b>Bootstrap Freeform Config (NXOS)</b>	<p>(Optional) Enter additional commands for NX-OS switches, as needed. For example, if you require some additional configurations to be pushed to the device and be available post device bootstrap, they can be captured in this field, to save the desired intent. After the devices boot up, they will contain the configuration defined in the <b>Bootstrap Freeform Config</b> field.</p> <p>Copy-paste the running-config to a <b>freeform config</b> field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. For more information, see <a href="#">Enabling Freeform Configurations on Fabric Switches</a>.</p>
<b>DHCPv4 Multi Subnet Scope</b>	<p>Specifies the field to enter one subnet scope per line. This field is editable after you check the <b>Enable Local DHCP Server</b> check box. The format of the scope should be defined as:</p> <p><b>DHCP Scope Start Address, DHCP Scope End Address, Switch Management Default Gateway, Switch Management Subnet Prefix</b></p> <p>For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24</p>

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

# Configuration Backup

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

Field	Description
<b>Hourly Fabric Backup</b>	Select the check box to enable an hourly backup of fabric configurations and the intent. The hourly backups are triggered during the first 10 minutes of the hour.
<b>Scheduled Fabric Backup</b>	Check the check box to enable a daily backup. This backup tracks changes in running configurations on the fabric devices that are not tracked by configuration compliance.
<b>Scheduled Time</b>	<p>Specify the scheduled backup time in a 24-hour format. This field is enabled if you check the <b>Scheduled Fabric Backup</b> check box.</p> <p>Select both the check boxes to enable both back up processes. The backup process is initiated after you click <b>Save</b>.</p> <p>The scheduled backups are triggered exactly at the time you specify with a delay of up to two minutes. The scheduled backups are triggered regardless of the configuration deployment status.</p> <p>The number of fabric backups that are retained on Nexus Dashboard is decided by the <b>Admin &gt; System Settings &gt; Server Settings &gt; LAN Fabric &gt; Maximum Backups per Fabric</b> option.</p> <p>The number of archived files that are retained is set in the <b># Number of archived files per device to be retained:</b> field on the <b>Server Properties</b> page.</p> <p>Note: To trigger an immediate backup, do the following:</p> <ol style="list-style-type: none"><li>1. Choose <b>Overview &gt; Topology</b>.</li><li>2. Click within the specific fabric box. The fabric topology screen comes up.</li><li>3. Right-click on a switch within the fabric, then select <b>Preview Config</b>.</li><li>4. On the <b>Preview Config</b> page for this fabric, click <b>Re-Sync All</b>.</li></ol> <p>You can also initiate the fabric backup on the fabric topology page. Click <b>Backup Now</b> in the <b>Actions</b> pane.</p>

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

## Border Gateway

The **Border Gateway** tab is applicable only to Cisco Nexus 9000 switches. The fields in the tab are described in the following table. Most of the fields are automatically populated based on Cisco-

recommended best practice configurations, but you can update the fields if needed.

Field	Description
<b>Site Id</b>	Specifies the ID for this fabric when you are moving this fabric within a VXLAN EVPN Multi-Site. The site ID is mandatory for a member fabric to be a part of a VXLAN EVPN Multi-Site. Each member fabric of a VXLAN EVPN Multi-Site has a unique site ID for identification.
<b>Anycast Border Gateway advertise-pip</b>	Advertises Anycast Border Gateway PIP as VTEP.
<b>Enable L3VNI w/o VLAN</b>	Enables the Layer 3 VNI without VLAN option. The setting at this fabric-level field affects the related field at the VRF level. For more information, see: <ul style="list-style-type: none"> <li>• <a href="#">Layer 3 VNI without VLAN</a></li> <li>• The "Creating a VRF" section in <a href="#">About Fabric Overview for LAN Operational Mode Setups</a></li> </ul>
<b>vPC Peer Link VLAN Range</b>	Specifies the VLAN range used for the vPC peer link SVI. The vPC fields become active only if the switch role is border gateway. Valid entries: 2-4094.
<b>Make vPC Peer Link VLAN as Native VLAN</b>	Enables vPC peer link VLAN as Native VLAN.
<b>vPC Peer Keep Alive option</b>	Allows you to configure routed links between vPC peers using management or loopback interfaces. <p>To use IP addresses assigned to the management port and the management VRF, choose <b>management</b>. To use IP addresses assigned to loopback interfaces and a non-management VRF, choose underlay routing loopback with IPv6 address for PKA. Both the options are supported for IPv6 underlay.</p>
<b>vPC Auto Recovery Time (In Seconds)</b>	Specifies the vPC auto recovery time-out period in seconds.
<b>vPC Delay Restore Time (In Seconds)</b>	Specifies the vPC delay restore period in seconds.
<b>vPC Peer Link Port Channel ID</b>	Specifies the Port Channel ID for a vPC Peer Link. By default, the value in this field is 500.
<b>vPC IPv6 ND Synchronize</b>	Enables IPv6 Neighbor Discovery synchronization between vPC switches. The check box is enabled by default.
<b>vPC advertise-pip</b>	Select the check box to enable the Advertise PIP feature. You can enable the advertise PIP feature also on a specific vPC.
<b>vPC Domain Id Range</b>	Specifies the vPC Domain Id range to use for new pairings.
<b>Enable NX-API</b>	Enables NX-API on HTTPS. This check box is checked by default.
<b>NX-API HTTPS Port Number</b>	Specifies the port on which NX-API is enabled. By default, NX-API is enabled on HTTPS port 443.

Field	Description
<b>Enable HTTP NX-API</b>	Enables NX-API to use HTTP connections. This option is enabled by default. However, it is recommended to use HTTPS for secure communication.
<b>NX-API HTTP Port Number</b>	Specifies the port on which NX-API is enabled. By default, NX-API is enabled on HTTP port 80.
<b>Enable TCAM Allocation</b>	Automatically generates TCAM commands for VXLAN and vPC Fabric Peering, when enabled.
<b>Nexus Border Gateway Freeform Config</b>	Allows you to configure additional CLIs for all the border gateway switches.
<b>Nexus Intra-fabric Links Additional Config</b>	Allows you to configure additional CLIs for all the intra fabric links.
<b>Greenfield Cleanup Option</b>	Enables cleaning up the switches imported into Nexus Dashboard with Preserve-Config=No, without a switch reload. This option is typically recommended only for the fabric environments with Cisco Nexus 9000v switches to improve the switch clean-up time.

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

# Telemetry

For a Nexus Dashboard fabric in Monitored mode, Nexus Dashboard will not deploy the telemetry configuration to the switches in the fabric. In order to activate, pause, or deactivate telemetry, you must copy and paste the expected configuration on the switches.



1. Navigate to **Admin > System Status > Telemetry > Switches**
2. Click the checkbox next to a switch, then click **Actions > Expected configuration**.
3. View and copy the configuration information in the **Software Telemetry** and **Flow Telemetry** area,.
4. Using the command line, log in to the switch.
5. Enter this command.

```
switch(config)# copy running-config startup-config
```

The telemetry feature in Nexus Dashboard allows you to collect, manage, and monitor real-time telemetry data from your Nexus Dashboard. This data provides valuable insights into the performance and health of your network infrastructure, enabling you to troubleshoot proactively and optimize operations. When you enable telemetry, you gain enhanced visibility into network operations and efficiently manage your fabrics.

Follow these steps to enable telemetry for a specific fabric.

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

Go to **Manage > Fabrics**.

2. Choose the fabric for which you want to enable telemetry.
3. From the **Actions** drop-down list, choose **Edit fabric settings**.

The **Edit *fabric-name* settings** page displays.



You can also access the **Edit *fabric-name* settings** page for a fabric from the **Fabric Overview** page. In the **Fabric Overview** page, click the **Actions** drop-down list and choose **Edit fabric settings**.

4. In the **Edit *fabric-name* settings** page, click the **General** tab.
5. Under the **Enabled features** section, check the **Telemetry** check box.
6. Click **Save**.

Navigate back to the **Edit *fabric-name* settings** page. The **Telemetry** tab displays.



The **Telemetry** tab appears only when you enable the **Telemetry** option under the **General** tab in the **Edit *fabric-name* settings** page.

The **Telemetry** tab includes these options.

- **Configuration**—allows you to manage telemetry settings and parameters.
- **NAS**—provides Network Analytics Service (NAS) features for advanced insights.

## Edit configuration settings

The **Configuration** tab includes these settings.

- **General**—allows you to enable analysis.

You can enable these settings.

- **Enable assurance analysis**—enables you to collect of telemetry data from devices to ensure network reliability and performance.
- **Enable Microburst sensitivity** - allows you to monitor traffic to detect unexpected data bursts within a very small time window (microseconds). Choose the sensitivity type from the **Microburst Sensitivity Level** drop-down list. The options are **High sensitivity**, **Medium sensitivity**, and **Low sensitivity**.



The **Enable Microburst sensitivity** option is available only for ACI fabrics.

- **Flow collection modes**—allows you to choose the mode for telemetry data collection. Modes include **NetFlow**, **sFlow**, and **Flow Telemetry**.

For more information see: [Flow collection](#) and [Configure flows](#).

- **Flow collection rules**—allows you to define rules for monitoring specific subnets or endpoints. These rules are pushed to the relevant devices, enabling detailed telemetry data collection.

For more information, see [Flow collection](#).

## Edit NAS settings

Nexus Dashboard allows you to export captured flow records to a remote NAS device using the Network File System (NFS) protocol. Nexus Dashboard defines the directory structure on NAS where the flow records are exported.

You can choose between the two export modes.

- **Full**—exports the complete data for each flow record.
- **Base**—exports only the essential 5-tuple data for each flow record.

Nexus Dashboard needs both read and write permissions on the NAS to perform the export successfully. If Nexus Dashboard cannot write to the NAS, it will generate an alert to notify you of the issue.

# Disable Telemetry

You can uncheck the **Telemetry** check box on your fabric's **Edit Fabric Settings > General >** page to disable the telemetry feature for your fabric. Disabling telemetry puts the telemetry feature in a transition phase and eventually the telemetry feature is disabled.

In certain situations, the disable telemetry workflow can fail, and you may see the **Force disable telemetry** option on your fabric's **Edit Fabric Settings** page.

If you disable the telemetry option using the instructions provided in [Perform force disable telemetry on your fabric](#), Nexus Dashboard acknowledges the user intent to disable telemetry feature for your fabric, ignoring any failures.

The Nexus Dashboard **Force disable telemetry** allows you to perform a force disable action for the telemetry configuration on your fabric. This action is recommended when the telemetry disable workflow has failed and you need to disable the telemetry feature on your fabric.



Using the **Force disable telemetry** feature may leave switches in your fabric with stale telemetry configurations. You must manually clean up these stale configurations on the switches before re-enabling telemetry on your fabric.

## Perform force disable telemetry on your fabric

Follow these steps to perform a force disable telemetry on your fabric.

1. (Optional) Before triggering a force disable of telemetry configuration, resolve any telemetry configuration anomalies flagged on the fabric.
2. On the **Edit Fabric Settings** page of your fabric, a banner appears to alert you that telemetry cannot be disabled gracefully, and a **Force Disable** option is provided with the alert message.
3. Disable telemetry from the Nexus Dashboard UI using one of these options.
  - a. Click the **Force disable** option in the banner that appears at the top of your fabric's **Edit Fabric Settings** page to disable telemetry for your fabric gracefully.
  - b. Navigate to your fabric's **Overview** page and click the **Actions** drop-down list to choose **Telemetry > Force disable telemetry** option.

Once the force disable action is executed, the **Telemetry** configuration appears as disabled in **Edit Fabric Settings > General > Enabled features > Telemetry** area, that is, the **Telemetry** check box is unchecked.

4. Clean up any stale telemetry configurations from the fabric before re-enabling telemetry on Nexus Dashboard.

## NAS

You can export flow records captured by Nexus Dashboard on a remote Network Attached Storage (NAS) with NFS.

Nexus Dashboard defines the directory structure on NAS where the flow records are exported.

You can export the flow records in Base or Full mode. In Base mode, only 5-tuple data for the flow record is exported. In Full mode the entire data for the flow record is exported.

Nexus Dashboard requires read and write permission to NAS in order to export the flow record. A system issue is raised if Nexus Dashboard fails to write to NAS.

## Guidelines and limitations for network attached storage

- In order for Nexus Dashboard to export the flow records to an external storage, the Network Attached Storage added to Nexus Dashboard must be exclusive for Nexus Dashboard.
- Network Attached Storage with Network File System (NFS) version 3 must be added to Nexus Dashboard.
- Flow Telemetry and Netflow records can be exported.
- Export of FTE is not supported.
- Average Network Attached Storage requirements for 2 years of data storage at 20k flows per sec:
  - Base Mode: 500 TB data
  - Full Mode: 2.8 PB data
- If there is not enough disk space, new records will not be exported and an anomaly is generated.

## Add network attached storage to export flow records

The workflow to add Network Attached Storage (NAS) to export flow records includes the following steps:

1. Add NAS to Nexus Dashboard.
2. Add the onboarded NAS to Nexus Dashboard to enable export of flow records.

### Add NAS to Nexus Dashboard

Follow these steps to add NAS to Nexus Dashboard.

1. Navigate to **Admin > System Settings > General**.
2. In the **Remote storage** area, click **Edit**.
3. Click **Add Remote Storage Locations**.
4. Complete the following fields to add NAS to Nexus Dashboard.
  - a. Enter the name of the Network Attached Storage and a description, if desired.
  - b. In the **Remote storage location type** field, click **NAS Storage**.
  - c. In the **Type** field, choose **Read Write**.

Nexus Dashboard requires read and write permission to export the flow record to NAS. A system issue is raised if Nexus Dashboard fails to write to NAS.

- d. In the **Hostname** field, enter the IP address of the Network Attached Storage.
- e. In the **Port** field, enter the port number of the Network Attached Storage.

f. In the **Export path** field, enter the export path.

Using the export path, Nexus Dashboard creates the directory structure in NAS for exporting the flow records.

g. In the **Alert threshold** field, enter the alert threshold time.

Alert threshold is used to send an alert when the NAS is used beyond a certain limit.

h. In the **Limit (Mi/Gi)** field, enter the storage limit in Mi/Gi.

i. Click **Save**.



In a three-node Nexus Dashboard cluster deployment, NAS backup operations require that all cluster nodes are in an operational and healthy state. If any node is down, unreachable, or not in a fully functional state, NAS backup operations do not proceed.

This ensures cluster consistency and data integrity during backup operations. Verify cluster health and confirm that all three nodes are up and synchronized before initiating or scheduling NAS backups.

## Add the onboarded NAS to Nexus Dashboard

Follow these steps to add the onboarded NAS to Nexus Dashboard.

1. Navigate to the Fabrics page:

**Manage > Fabrics**

2. Choose the fabric with the telemetry feature enabled.

3. Choose **Actions > Edit Fabric Settings**.

4. Click **Telemetry**.

5. Click the **NAS** tab in the **Telemetry** window.

6. Make the necessary configurations in the **General settings** area.

a. Enter the name in the **Name** field.

b. In the **NAS server** field, choose the NAS server added to Nexus Dashboard from the drop-down list.

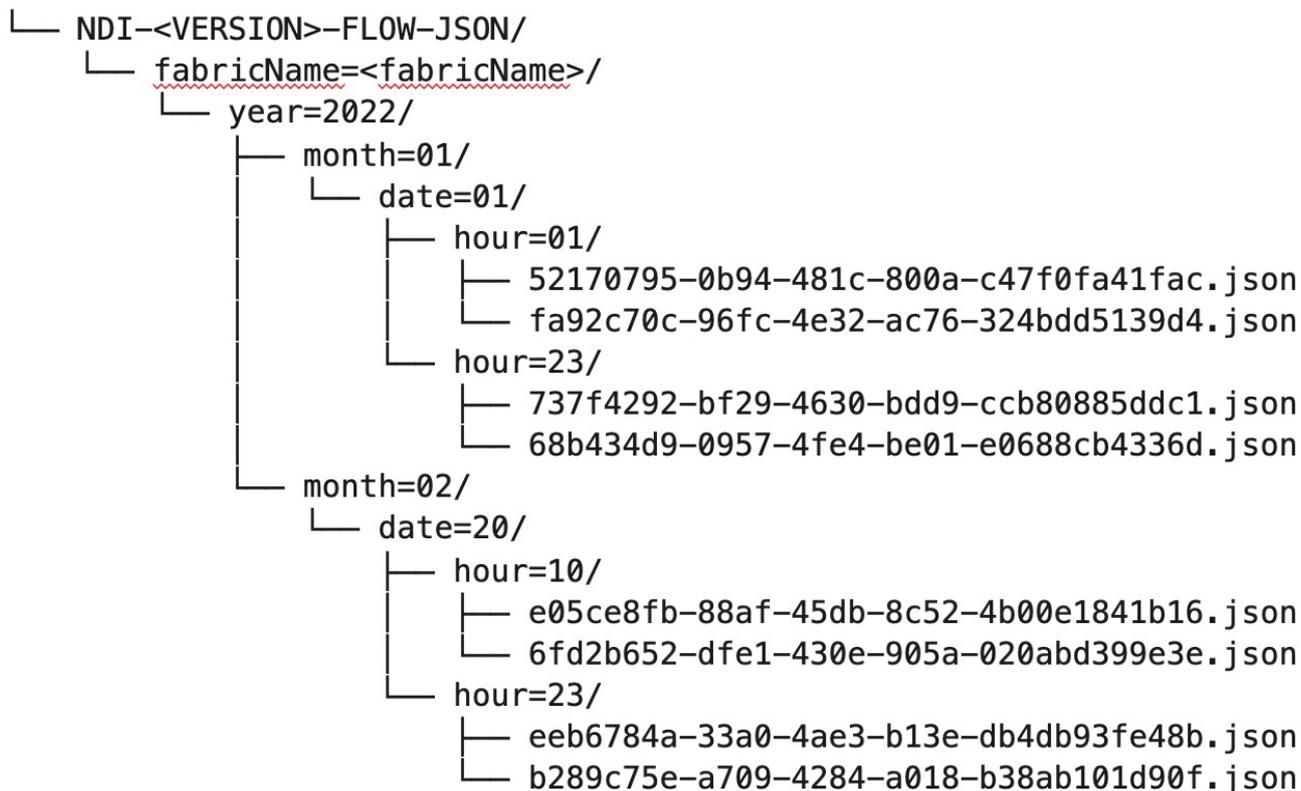
7. In the **Collection settings** area, choose the flow from the **Flows** drop-down list.

o In Base mode, only 5-tuple data for the flow record is exported.

o In Full mode, the entire data for the flow record is exported.

8. Click **Save**.

The traffic from the flows displayed in the **Flows** page is exported as a JSON file to the external NAS in the following directory hierarchy.



Navigate to **Analyze > Flows** to view the flows that will be exported.

Each flow record is written as a line delimited JSON.

#### JSON output file format for a flow record in base mode

```

{"fabricName":"myopic","terminalTs":1688537547433,"originTs":1688537530376,"srcIpl":
"2000:201:1:1::1","dstIpl":"2000:201:1:1::3","srcPort":1231,"dstPort":1232,"ingressVrf":
"vrf1","egressVrf":"vrf1","ingressTenant":"FSV1","egressTenant":"FSV1","protocol":"U
DP"}

```

```

{"fabricName":"myopic","terminalTs":1688537547378,"originTs":1688537530377,"srcIpl":
"201.1.1.127","dstIpl":"201.1.1.1","srcPort":0,"dstPort":0,"ingressVrf":"vrf1","egressVrf":
":"","ingressTenant":"FSV2","egressTenant":"","protocol":"ANY-HOST"}

```

#### JSON output file format for a flow record in full mode

```

{"fabricName":"myopic","terminalTs":1688538023562,"originTs":1688538010527,"srcIpl":
"201.1.1.121","dstIpl":"201.1.1.127","srcPort":0,"dstPort":0,"ingressVrf":"vrf1","egress
Vrf":"vrf1","ingressTenant":"FSV2","egressTenant":"FSV2","protocol":"ANY-
HOST","srcEpg":"ext-epg","dstEpg":"ext-
epg1","latencyMax":0,"ingressVif":"eth1/15","ingressVni":0,"latency":0,"ingressNodes":
"Leaf1-
2","ingressVlan":0,"ingressByteCount":104681600,"ingressPktCount":817825,"ingressBur
st":0,"ingressBurstMax":34768,"egressNodes":"Leaf1-2","egressVif":"po4",

```

```
" egressVni" :0," egressVlan" :0," egressByteCount" :104681600," egressPktCount" :817825,"
 egressBurst" :0," egressBurstMax" :34768," dropPktCount" :0," dropByteCount" :0," dropCode
 ":" " , " dropScore" :0," moveScore" :0," latencyScore" :0," burstScore" :0," anomalyScore" :0,"
 hashCollision" :false," dropNodes" : " [ ]" , " nodeNames" : " [ \" Leaf1-
 2\ " ]" , " nodeIngressVifs" : " [ \" Leaf1-2,eth1/15\ " ]" , " nodeEgressVifs" : " [ \" Leaf1-2,po4\ " ]“
 , " srcMoveCount" :0," dstMoveCount" :0," moveCount" :0," prexmit" :0," rtoOutside" :false," ev
 ents" : " [[\\\" 1688538010527,Leaf1-2,0,3,1,no,no,eth1/15,,po4,po4,,,,,0,64,0,,,,,,\\\" ]]" }
```

## Flow collection

### Understanding flow telemetry

Flow telemetry allows users to see the path taken by different flows in detail. It also allows you to identify the EPG and VRF instance of the source and destination. You can see the switches in the flow with the help of flow table exports from the nodes. The flow path is generated by stitching together all the exports in order of the flow.

You can configure the Flow Telemetry rule for the following interface types:

- VRF instances
- Physical interfaces
- Port channel interfaces
- Routed sub-interfaces (Cisco ACI fabric)
- SVIs (Cisco ACI fabric)



In a Cisco ACI fabric, if you want to configure routed sub-interfaces from the UI, select L3 Out.

In an NX-OS fabric, physical or port channel flow rules are supported only on routed interfaces.

Flow telemetry monitors the flow for each fabric separately, as there is no stitching across the fabrics in a fabric group. Therefore, flow telemetry is for individual flows. For example, if there are two fabrics (fabric A and fabric B) within a fabric group, and traffic is flowing between the two fabrics, they will be displayed as two separate flows. One flow will originate from Fabric A and display where the flow exits. And the other flow from Fabric B will display where it enters and where it exits.

### Flow telemetry guidelines and limitations

- All flows are monitored as a consolidated view in a unified pipeline for Cisco ACI and NX-OS fabrics, and the flows are aggregated under the same umbrella.
- Even if a particular node (for example, a third-party switch) is not supported for Flow Telemetry, Nexus Dashboard will use LLDP information from the previous and next nodes in the path to identify the switch name and the ingress and egress interfaces.
- Nexus Dashboard supports Kafka export for Flow anomalies. However, Kafka export is not currently supported for Flow Event anomalies.

## Flow telemetry guidelines and limitations for NX-OS fabrics

- Ensure that you have configured NTP and enabled PTP in Nexus Dashboard. See [Cisco Nexus Dashboard Deployment Guide](#) and [Precision Time Protocol \(PTP\) for Cisco Nexus Dashboard Insights](#) for more information. You are responsible for configuring the switches with external NTP servers.
- In the **Edit Flow** page, you can enable all three telemetry types. sFlow is most restrictive, Netflow has some more capability, and Flow Telemetry has the most capability. We recommend that you enable Flow Telemetry if it is available for your configuration. If Flow Telemetry is not available, then use Netflow. If Netflow is not available, use sFlow.
- If there are multiple Nexus Dashboard clusters onboarded to Nexus Dashboard, partial paths will be generated for each fabric.
- If you manually configure the fabric to use with Nexus Dashboard and Flow Telemetry support, the Flows Exporter port changes from 30000 to 5640. To prevent a breakage of the Flows Exporter, adjust the automation.
- Nexus Dashboard supports Kafka export for Flow anomalies. However, Kafka export is not currently supported for Flow Event anomalies.
- Flow telemetry is supported in -FX3 platform switches for the following NX-OS versions:
  - 9.3(7) and later
  - 10.1(2) and later
  - Flow telemetry is not supported in -FX3 platform switches for NX-OS version 10.1(1).
- Interface based Flow Telemetry is only supported on modular chassis with -FX land -GX line cards on physical ports and port-channels rules.
- If interface-based Flow Telemetry is pushed from Nexus Dashboard for **Classic LAN** and **External Connectivity Network** fabrics, perform the following steps:
  - Choose the fabric.
  - Choose **Policies** > **Action** > **Add policy** > **Select all** > **Choose template** > host\_port\_resync and click **Save**.
  - In the Fabric Overview page, choose **Actions** > **Recalculate and deploy**.
- For VXLAN fabrics, interface-based Flow Telemetry is not supported on switch links between spine switch and leaf switch.
- If you want to use the default VRF instance for flow telemetry, you must create the VRF instance with a name of "default" in lowercase. Do not enter the name with any capital letters.
- Flow telemetry is not supported in classic LAN topologies with 2-level VPC access layers.
- If you want to enable Flow Telemetry, ensure that there are no pre-existing Netflow configurations on the switches. If there are any pre-existing configurations, the switch configuration may fail.

To enable Flow Telemetry without configuration issues, follow these steps:

- Ensure that there are no pre-existing Netflow configurations on the switches. If such configurations exist, enabling Flow Telemetry might result in a system anomaly with an error message stating **invalid command match IP source address**.
- If you encounter the error, disable Flow Telemetry.
- Remove any existing Netflow configurations from the switches.

- o Re-enable Flow Telemetry.
- o For some flows, latency information is not available, which could happen due to latency issues. In these cases, latency information will be reported as 0.

## Flow telemetry rules guidelines and limitations for NX-OS fabrics

- If you configure an interface rule (physical or port channel) on a subnet, it can monitor only incoming traffic. It cannot monitor outgoing traffic on the configured interface rule.
- If a configured port channel that contains two physical ports, only the port channel rule is applicable. Even if you configure physical interface rules on the port, only port channel rule takes precedence.
- For NX-OS release 10.3(2) and earlier, if a flow rule are configured on an interface, then global flow rules are not matched.
- For NX-OS release 10.3(3) and later, a flow rule configured on an interface is matched first and then the global flow rules are matched.

## Configure flows

### Configure flow collection modes

Follow these steps to configure flow collection modes.

1. Navigate to **Admin > System Settings > Flow collection**.
2. In the **Flow collection mode** area, choose **Flow telemetry**.



Enabling Flow Telemetry automatically activates Flow Telemetry Events. Whenever a compatible event takes place, an anomaly will be generated, and the What's the impact? section in the **Anomaly** page will display the associated flows. You can manually configure a Flow Telemetry rule to acquire comprehensive end-to-end information about the troublesome flow.

### Configure flow collection rules in an NX-OS fabric

Follow these steps to configure flow collection rules in an NX-OS fabric.

1. Navigate to the **Telemetry** window for your fabric.
  - a. Navigate to the main **Fabrics** page:  
**Manage > Fabrics**
  - b. In the table showing all of the Nexus Dashboard fabrics that you have already created, locate the LAN or IPFM fabric where you want to configure telemetry settings.
  - c. Single-click on that fabric.  
  
The **Overview** page for that fabric appears.
  - d. Click **Actions > Edit Fabric Settings**.

The **Edit *fabric\_name* Settings** window appears.

e. Verify that the **Telemetry** option is enabled in the **Enabled features** area.

The Telemetry tab doesn't become available unless the **Telemetry** option is enabled in the **Enabled features** area.

f. Click the **Telemetry** tab to access the telemetry settings for this fabric.

2. Click the **Flow collection** tab in the **Telemetry** window.

3. In the **Mode** area, click **Flow telemetry**.

4. In the **Flow collections rules** area, determine what sort of flow collection rule that you want to add.

- o [VRF](#)
- o [Physical interface](#)
- o [Port channel](#)

## VRF

To add a VRF rule:

1. Click the **VRF** tab.

A table with already-configured VRF flow collection rules is displayed.

For any VRF flow collection rule in this table, click the ellipsis (...), then click **Edit rule** to edit that rule or **Delete rule** to delete it.

2. Add a new rule by clicking **Create flow collection rule**.

a. In the **General** area, complete the following:

- i. Enter the name of the rule in the **Rule Name** field.
- ii. The VRF field is disabled. The flow rule applies to all the VRF instances.
- iii. In the **Flow Properties** area, select the protocol for which you intend to monitor the flow traffic.
- iv. Enter the source and destination IP addresses. Enter the source and destination port.
- v. Click **Save**.

## Physical interface

To add a physical interface rule:

1. Click the **Physical interface** tab.

A table with already-configured physical interface flow collection rules is displayed.

For any physical interface flow collection rule in this table, click the ellipsis (...), then click **Edit rule** to edit that rule or **Delete rule** to delete it.

2. Add a new rule by clicking **Create flow collection rule**.

a. In the **General** area, complete the following:

- i. Enter the name of the rule in the **Rule Name** field.

- ii. Check the **Enabled** check box to enable the status. If you enable the status, the rule will take effect. Otherwise, the rule will be removed from the switches.
- iii. In the **Flow Properties** area, select the protocol for which you intend to monitor the flow traffic.
- iv. Enter the source and destination IP addresses. Enter the source and destination port.
- v. In the **Interface List** area, click **Select a Node**. Use the search box to select a node.
- vi. From the drop-down list, select an interface. You can add more than one row (node+interface combination) by clicking **Add Interfaces**. However, within the rule, a node can appear only once. Configuration is rejected if more than one node is added.
- vii. Click **Save**.

### Port channel

To add a port channel rule:

1. Click the **Port channel** tab.

A table with already-configured port channel flow collection rules is displayed.

For any port channel flow collection rule in this table, click the ellipsis (...), then click **Edit rule** to edit that rule or **Delete rule** to delete it.

2. Add a new rule by clicking **Create flow collection rule**.

- a. In the **General** area, enter the name of the rule in the **Rule Name** field.

- i. Select the **Enabled** check box to enable the status. If you enable the status, the rule will take effect. Otherwise, the rule will be removed from the switches.
- ii. In the **Flow Properties** area, select the protocol for which you intend to monitor the flow traffic.
- iii. Enter the source and destination IP addresses. Enter the source and destination port.
- iv. From the drop-down list, select an interface. You can add more than one row (node+interface combination) by clicking **Add Interfaces**. However, within the rule, a node can appear only once. Configuration is rejected if more than one node is added.
- v. Click **Save**.

3. Click **Done**.

### Monitor the subnet for flow telemetry

In the following example, the configured rule for a flow monitors the specific subnet provided. The rule is pushed to the fabric which pushes it to the switches. So, when the switch sees traffic coming from a source IP or the destination IP, and if it matches the subnet, the information is captured in the TCAM and exported to the Nexus Dashboard service. If there are 4 nodes (A, B, C, D), and the traffic moves from A > B > C > D, the rules are enabled on all 4 nodes and the information is captured by all the 4 nodes. Nexus Dashboard stitches the flows together. Data such as the number of drops and the number of packets, anomalies in the flow, and the flow path are aggregated for the 4 nodes.

Follow these steps to monitor the subnet for flow telemetry.

1. Navigate to **Manage > Fabric**.
2. Choose a fabric.
3. Verify that your **Fabrics** and the **Snapshot** values are appropriate. The default snapshot value is 15 minutes. Your choice will monitor all the flows in the chosen fabric or snapshot fabric.
4. Navigate to **Connectivity > Flows** to view a summary of all the flows that are being captured based on the snapshot that you chose.

The related anomaly score, record time, the nodes sending the flow telemetry, flow type, ingress and egress nodes, and additional details are displayed in a table format. If you click a specific flow in the table, specific details are displayed in the sidebar for the particular flow telemetry. In the sidebar, if you click the Details icon, the details are displayed in a larger page. In this page, in addition to other details, the **Path Summary** is also displayed with specifics related to source and destination. If there are flows in the reverse direction, that will also be visible in this location.

For a bi-directional flow, there is an option to choose to reverse the flow and see the path summary displayed. If there are any packet drops that generate a flow event, they can be viewed in the Anomaly dashboard.

# Understanding Netflow

Netflow is an industry standard where Cisco routers monitor and collect network traffic on an interface. Netflow version 9 is supported.

Netflow enables the network administrator to determine information such as source, destination, class of service, and causes of congestion. Netflow is configured on the interface to monitor every packet on the interface and provide telemetry data. You cannot filter on Netflow.

Netflow in Nexus series switches is based on intercepting the packet processing pipeline to capture summary information of network traffic.

The components of a flow monitoring setup are as follows:

- Exporter: Aggregates packets into flows and exports flow records towards one or more collectors
- Collector: Reception, storage, and pre-processing of flow data received from a flow exporter
- Analysis: Used for traffic profiling or network intrusion
- The following interfaces are supported for Netflow:

## Supported interfaces for Netflow

Interfaces	5 Tuple	Nodes	Ingress	Egress	Path	Comments
Routed Interface/Port Channel	Yes	Yes	Yes	No	Yes	Ingress node is shown in path
Sub Interface/Logical (Switch Virtual Interface)	Yes	Yes	No	No	No	No



In an NX-OS fabric, port channel support is available if you monitor only the host-facing interfaces.

## Understanding Netflow types

You can use these Netflow types.

- [Full Netflow](#)
- [Sampled Netflow](#)

### Full Netflow

With Full Netflow, all packets on the configured interfaces are captured into flow records in a flow table. Flows are sent to the supervisor module. Records are aggregated over configurable intervals and exported to the collector. Except in the case of aliasing (multiple flows hashing to the same entry in the flow table), all flows can be monitored regardless of their packet rate.

Nexus 9000 Series switches with the Fabric Controller type as well as switches in a Cisco ACI fabric

support Full Netflow.

## Sampled Netflow

With Sampled Netflow, packets on configured interfaces are time sampled. Flows are sent to the supervisor or a network processor for aggregation. Aggregated flow records are exported at configured intervals. The probability of a record for a flow being captured depends on the sampling frequency and packet rate of the flow relative to other flows on the same interface.

Nexus 7000 and Nexus 7700 Series switches with F/M line cards and the Fabric Controller type, support Sampled Netflow.

## Netflow guidelines and limitations

- In Cisco Nexus 9000 series switches, Netflow supports a small subset of the published export fields in the RFC.
- Netflow is captured only on the ingress port of a flow as only the ingress switch exports the flow. Netflow cannot be captured on fabric ports.

### Netflow guidelines and limitations for Cisco ACI fabrics

- We recommend that you enable Flow Telemetry. If that is not available for your configuration, use Netflow. However, you can determine which mode of flow to use based upon your fabric configuration.
- Enabling both Flow Telemetry and Netflow is not supported.
- After you enable Netflow, you must obtain the Netflow collector IP address and configure Cisco APIC with the collector IP address. See [Cisco APIC and NetFlow](#).

To obtain the Netflow collector IP address, navigate to **Admin > System Settings > General**, then locate the **External Pools** area. Click **View all** at the bottom left area of the **External Pools** tile; the **Telemetry-collector** persistent IP addresses listed in the table are used for the Netflow collector IP address.

- The Netflow and sFlow flow collection modes do not support any anomaly.

### Netflow guidelines and limitations for NX-OS fabrics

- In the **Edit Flow** page, you can enable all three modes. Choose the best possible mode for a product. sFlow is the most restrictive, Netflow has more capabilities, and Flow Telemetry has the most capabilities. We recommend that you enable Flow Telemetry if it is available for your configuration. If Flow Telemetry is not available, then use Netflow. If Netflow is not available, use sFlow.
- In Nexus 7000 and Nexus 9000 Series switches, only the ingress host-facing interface configured for Netflow are supported (either in VXLAN or Classic LAN).
- The Netflow supported fabrics are Classic and VXLAN. VXLAN is not supported on fabric ports.
- Netflow configurations will not be pushed. However, if a fabric is managed, the software sensors will be pushed.
- If you manually configure the fabric to use with Nexus Dashboard and Netflow support, the Flows Exporter port changes from 30000 to 5640. To prevent a breakage of the Flows Exporter, adjust

the automation.

- To configure Netflow on fabric switches, see the **Configuring Netflow** section in the [Cisco Nexus 9000 Series NX-OS System Management Configuration Guide](#).

## Configure Netflow

Follow these steps to configure Netflow.

1. Navigate to the **Telemetry** page for your LAN or IPFM fabric.
2. Click the **Flow collection** tab on the **Telemetry** page.
3. In the **Mode** area, make the following choices:
  - Choose **Netflow**.
  - Choose **Flow Telemetry**.
4. Click **Save**.

# Understanding sFlow

sFlow is an industry standard technology traffic in data networks containing switches and routers. Nexus Dashboard supports [sFlow version 5](#) on Cisco Nexus 3000 series switches.

sFlow provides the visibility to enable performance optimization, an accounting and billing for usage, and defense against security threats.

The following interfaces are supported for sFlow:

## *Supported interfaces for sFlow*

Interfaces	5 Tuple	Nodes	Ingress	Egress	Path	Comments
Routed Interface	Yes	Yes	Yes	Yes	Yes	Ingress node is shown in path

## Guidelines and limitations for sFlow

- Nexus Dashboard supports sFlow with Cisco Nexus 3000 series switches.
- It is recommended to enable Flow Telemetry if it is available for your configuration. If it is not available for your configuration, use Netflow. If Netflow, is not available for your configuration, then use sFlow.
- For sFlow, Nexus Dashboard requires the configuration of persistent IPs under cluster configuration, and 6 IPs in the same subnet as the data network are required.
- sFlow configurations will not be pushed. However, if a fabric is managed, the software sensors will be pushed.
- If you manually configure the fabric to use with Nexus Dashboard and sFlow support, the Flows Exporter port changes from 30000 to 5640. To prevent a breakage of the Flows Exporter, adjust the automation.
- Nexus Dashboard does not support sFlow in the following Cisco Nexus 3000 Series switches:
  - Cisco Nexus 3600-R Platform Switch (N3K-C3636C-R)
  - Cisco Nexus 3600-R Platform Switch (N3K-C36180YC-R)
  - Cisco Nexus 3100 Platform Switch (N3K-C3132C-Z)
- Nexus Dashboard does not support sFlow in the following Cisco Nexus 9000 Series fabric modules:
  - Cisco Nexus 9508-R fabric module (N9K-C9508-FM-R)
  - Cisco Nexus 9504-R fabric module (N9K-C9504-FM-R)
- To configure sFlow on fabric switches, see the **Configuring sFlow** section in the [Cisco Nexus 9000 Series NX-OS System Management Configuration Guide](#).

## Configure sFlow telemetry

### Prerequisites

Follow these steps to configure sFlow telemetry.

1. Navigate to the **Telemetry** page for your LAN or IPFM fabric.
2. Click the **Flow collection** tab on the **Telemetry** page.
3. In the **Mode** area, make the following choices:
  - o Choose **sFlow**.
  - o Choose **Flow Telemetry**.
4. Click **Save**.

# External streaming

The **External streaming** tab in Nexus Dashboard allows you export data that Nexus Dashboard collects over Kafka, email, and syslog. Nexus Dashboard generates data such as advisories, anomalies, audit logs, faults, statistical data, and risk and conformance reports. When you configure a Kafka broker, Nexus Dashboard writes all data to a topic. By default, the Nexus Dashboard collects export data every 30 seconds or at a less frequent interval.

For ACI fabrics, you can also collect data for specific resources (CPU, memory, and interface utilization) every 10 seconds from the leaf and spine switches using a separate data pipeline. To export this data, select the **Usage** option under **Collection Type** in the **Message bus** export settings. Additionally, CPU and memory data is collected for the controllers.



Nexus Dashboard does not store the collected data in Elasticsearch; instead, it exports the data directly to your repository or data lake using a Kafka broker for consumption. By using the Kafka export functionality, you can then export this data to your Kafka broker and push it into your data lake for further use.

You can configure an email scheduler to define the type of data and the frequency at which you want to receive information via email. You can also export anomaly records to an external syslog server. To do this, select the **Syslog** option under the **External Streaming** tab.

## Configure external streaming settings

Follow these steps to configure external streaming settings.

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

Go to **Manage > Fabrics**.

2. Choose the fabric for which you configure streaming settings.
3. From the **Actions** drop-down list, choose **Edit fabric settings**.

The **Edit *fabric-name* settings** page displays.



You can also access the **Edit *fabric-name* settings** page for a fabric from the **Fabric Overview** page. In the **Fabric Overview** page, click the **Actions** drop-down list and choose **Edit fabric settings**.

4. In the **Edit *fabric-name* settings** page, click the **External streaming** tab.

You can view these options.

- o Email
- o Message bus
- o Syslog

## Guidelines and limitations

- Intersight connectivity is required to receive the reports by email.
- You can configure up to five emails per day for periodic job configurations.
- A maximum of six exporters is supported for export across all types of exporters including email, message bus, and syslog. You must provide unique names for each export.
- The scale for Kafka exports is increased to support up to 20 exporters per cluster. However, statistics selection is limited to any six exporters.
- Before configuring your Kafka export, you must add the external Kafka IP address as a known route in your Nexus Dashboard cluster configuration and verify that Nexus Dashboard can reach the external Kafka IP address over the network.
- The anomalies in Kafka and email messages include categories such as Resources, Environmental, Statistics, Endpoints, Flows, and Bugs.
- Export data is not supported for snapshot fabrics.
- You must provide unique names for each exporter, and they may not be repeated between Kafka export for **Alerts and Events** and Kafka export for **Usage**.
- Nexus Dashboard supports Kafka export for flow anomalies. However, Kafka export is not currently supported for flow Event anomalies.

### Guidelines and limitations in NX-OS fabrics

- Remove all configurations in the *Message Bus Configuration* and *Email* page before you disable Software Telemetry on any fabric and remove the fabric from Nexus Dashboard.
- If you use interface groups while enabling PTP in an Enhanced Classic LAN fabric, you must add the following commands to the freeform section of the interface group.
  - `ttag`
  - `ttag-strip`

If you do not add these commands, endpoints connected behind these interfaces will experience traffic loss.

## Email

The email scheduler feature in Nexus Dashboard automates the distribution of summarized data collected from Nexus Dashboard. It allows customization of selection of email recipients, choice of email format, scheduling frequency settings, and configuring the types of alerts and reports.

Follow these steps to configure an email scheduler.

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

Go to **Manage > Fabrics**.

2. Choose the fabric for which you configure streaming settings.
3. From the **Actions** drop-down list, choose **Edit fabric settings**.

The **Edit *fabric-name* settings** page displays.

4. In the **Edit *fabric-name* settings** page, click the **External streaming** tab.
5. Click the **Email** tab.
6. Review the information provided in the **Email** tab for already-configured email configurations.

The following details display under **Email** tab.

Field	Description
<b>Name</b>	The name of the email configuration.
<b>Email</b>	The email addresses used in the email configuration.
<b>Start time</b>	The start date used in the email configuration.
<b>Frequency</b>	The frequency in days or weeks set in the email configuration.
<b>Anomalies</b>	The severity level for anomalies and advisories set in the email configuration.
<b>Advisories</b>	
<b>Risk and conformance reports</b>	The status of the overall inventory for a fabric, including software release, hardware platform, and a combination of software and hardware conformance.

To add a new email configuration, click **Add email** in the **Email** page.

1. Follow these steps to configure **General Settings**.
  - a. In the **Name** field, enter the name of the email scheduler.
  - b. In the **Email** field, enter one or more email addresses separated by commas.
  - c. In the **Format** field, choose **Text** or **HTML** email format.
  - d. In the **Start date** field, choose the start date when the scheduler should begin sending emails.
  - e. In the **Collection frequency in days** field, specify how often the summary is sent, you can choose days or weeks.

**General settings**

Name \*

Email \*

Use commas to enter multiple email addresses

**Format**

Text  HTML

**Start date\***

**Collection frequency in days\***

 Days ▾

2. Follow these steps to configure **Collection Settings**.
  - a. In the **Mode** field, choose one of the following modes.

- **Basic**—displays the severity levels for anomalies and advisories.
  - **Advanced**—displays the categories and severity levels for anomalies and advisories.
- b. Check the **Only include active alerts in email** check box, to include only active anomaly alerts.
  - c. Under **Anomalies** choose the categories and severity levels for the anomalies.
  - d. Under **Advisories** choose the categories and severity levels for the advisories.
  - e. Under **Risk and Conformance Reports**, choose from the following options.
    - Software
    - Hardware

Collection settings

Mode  
 Basic  Advanced

Only include active alerts in email

Anomalies [Select all](#) [Clear all](#)

Critical  
 Major  
 Warning  
 Minor

Advisories [Select all](#) [Clear all](#)

Critical  
 Major  
 Warning  
 Minor

Risk and Conformance Reports [Select all](#) [Clear all](#)

Software  
 Hardware

[Cancel](#) [Save](#)

3. Click **Save**.

The **Email** area displays the configured email schedulers.

You will receive an email about the scheduled job on the provided *Start Date* and at the time provided in the **Collection frequency in days** field. The subsequent emails follow after *Collect Every* frequency expires. If the provided time is in the past, Nexus Dashboard will send you an email immediately and trigger the next email after the duration from the provided start time expires.

## Message bus

### Add Kafka broker configuration

Follow these steps to configure the message bus and add kafka broker.

1. Configure the message bus at the **System Settings** level.
  - a. Navigate to **Admin > System Settings > General**.
  - b. In the **Message bus configuration** area, click **Edit**.

The **Message bus configuration** dialog box opens.

- c. Click **Add message bus configuration**.

The **Add message bus configuration** dialog box opens.

- d. In the **Name** field, enter a name for the configuration.
- e. In the **Hostname/IP address** and **Port** fields, enter the IP address of the message bus consumer and the port that is listening on the message bus consumer.
- f. In the **Topic name** field, enter the name of the Kafka topic to which Nexus Dashboard must send the messages.
- g. In the **Mode** field, choose the security mode.

The supported modes are **Unsecured**, **Secured SSL** and **SASLPLAIN**. The default value is **Unsecured**.

- For **Unsecured**, no other configurations are needed.
- For **Secured SSL**, fill out the following field:

**Client certification name**—The System Certificate name configured at the Certificate Management level. The CA certificate and System Certificate (which includes Client certificate and Client key) are added at the Certificate Management level.

Refer to Step 2 for step-by-step instructions on managing certificates. Navigate to **Admin > Certificate Management** to manage the following certificates:

- **CA Certificate**—The CA certificate used for signing consumer certificate, which will be stored in the trust-store so that Nexus Dashboard can trust the consumer.
  - **Client Certificate**—The CA signed certificate for Nexus Dashboard. The certificate is signed by the same CA, and the same CA certificate will be in the truststore of the consumer. This will be stored in Nexus Dashboard's Kafka keystore that is used for exporting.
  - **Client Key**—A private key for the Kafka producer, which is Nexus Dashboard in this case. This will be stored in Nexus Dashboard's Kafka keystore that is used for exporting.
- For **SASLPLAIN**, fill out these fields:
    - **Username**—The username for the SASL/PLAIN authentication.
    - **Password**—The password for the SASL/PLAIN authentication.

- h. Click **Save**

2. Add CA certificates and System certificates at the **Certificate Management level**.

- a. Navigate to **Admin > Certificate Management**.
- b. In the **Certificate management** page, click the **CA Certificates** tab, then click **Add CA certificate**.

The fields in the **CA Certificates** tab are described in the following table.

Field	Description
<b>Certificate name</b>	The name of the CA certificate.
<b>Certificate details</b>	The details of the CA certificate.
<b>Attached to</b>	The CA signed certificate attached to Nexus Dashboard.
<b>Expires on</b>	The Expiry date and time of the CA certificate.
<b>Last updated time</b>	The last updated time of the CA certificate.

- c. In the **Certificate management** page, click the **System certificates** tab, then click **Add system certificate** to add Client Certificate and Client key. Note that the Client certificate and Client key should have same names except extensions as .cer/.crt/.pem for Client certificate and .key for Client key.



You must add a valid CA Certificate before adding the corresponding System Certificate.

The fields in the **System Certificates** tab are described in the following table.

Field	Description
<b>Certificate name</b>	The name of the Client certificate.
<b>Certificate details</b>	The details of the Client certificate.
<b>Attached to</b>	The feature to which the system certificate is attached to, in this case, the message bus.
<b>Expires on</b>	The Expiry date and time of the CA certificate.
<b>Last updated time</b>	The last updated time of the CA certificate.



To configure message bus, the System Certificate should be attached to message bus feature.

To attach a System Certificate to the message bus feature:

- a. Choose the System Certificate that you want to use and click the ellipses (...) on that row.
- b. Choose **Manage Feature Attachments** from the drop-down list.

The **Manage Feature Attachments** dialog box opens.

- c. In the **Features** field, choose **messageBus**.
- d. Click **Save**.

For more information on CA certificates, see [Managing Certificates in your Nexus Dashboard](#).

## Configure Kafka exports in fabric settings

1. Navigate to the **External streaming** page for your fabric.
  - a. Navigate to the **Fabrics** page.

Go to **Manage > Fabrics**.

- b. Choose the fabric for which you configure streaming settings.
- c. From the **Actions** drop-down list, choose **Edit fabric settings**.

The **Edit fabric-name settings** page displays.

- d. In the **Edit fabric-name settings** page, click the **External streaming** tab.
  - e. Click the **Message bus** tab.
2. Review the information provided in the **Message bus** tab for already-configured message bus configurations, or click **Add message bus** to add a new message bus configuration.

Skip to Step 3 if you are adding a message bus.

The fields in the **Message bus** tab are described in the following table.

Field	Description
<b>Message bus stream</b>	The name of the message bus stream configuration.
<b>Collection type</b>	The collection type used by the message bus stream.
<b>Mode</b>	The mode used by the message bus stream.
<b>Anomalies</b>	The severity level for anomalies and advisories set in the message bus stream configuration.
<b>Advisories</b>	
<b>Statistics</b>	The statistics that were configured for the message bus stream.
<b>Faults</b>	The severity level for faults set in the message bus stream configuration.
<b>Audit Logs</b>	The audit logs that were configured for the message bus stream.

- 3. To configure a new message bus stream, in the **Message bus** page, click **Add message bus**.
- 4. In the **Message bus stream** field, choose the message bus stream that you want to edit.
- 5. In the **Collection Type** area, choose the appropriate collection type.

Depending on the **Collection Type** that you choose, the options displayed in this area will change.

- o **Alerts and events:** This is the default setting. Continue to Step 7, if you choose **Alerts and events**.
- o **Usage:** In the **Collection settings** area, under **Data**, the Resources, and Statistics for the collection settings are displayed. By default, the data for CPU, Memory, and Interface Utilization are collected and exported. You cannot choose to export a subset of these resources.



Usage is applicable only for ACI Fabrics. This option is disabled for other fabrics.

- 6. Click **Save**. The configured message bus streams are displayed in the **Message bus** area. This configuration now sends immediate notification when the selected anomalies or advisories occur.
- 7. If you choose **Alerts and events** as the **Collection Type**, in the **Mode** area, choose either **Basic** or **Advanced**.

The configurations that are available in each collection settings section might vary, depending on

the mode that you set.

8. Determine which area you want to configure for the message bus stream.

The following areas appear in the page:

- [Anomalies](#)
- [Advisories](#)
- [Statistics](#)
- [Faults](#)
- [Audit Logs](#)

After you complete the configurations on this page, click **Save**. Nexus Dashboard displays the configured message bus streams in the **Message bus** area. This configuration now sends immediate notification when the selected anomalies or advisories occur.

## Anomalies

- If you chose **Basic** in the **Mode** area, choose one or more of the following severity levels for anomaly statistics that you want to configure for the message bus stream:

- Critical
- Major
- Warning
- Minor

Or click **Select all** to select all available statistics for the message bus stream.

- If you chose **Advanced** in the **Mode** area:

- Choose one or more of the following categories for anomaly statistics that you want to configure for the message bus stream:

- Active Bugs
- Capacity
- Compliance
- Configuration
- Connectivity
- Hardware
- Integrations
- System

- Choose one or more of the following severity levels for anomaly statistics that you want to configure for the message bus stream:

- Critical
- Major
- Warning

- Minor

Or click **Select all** to select all available categories and statistics for the message bus stream. For more information on anomaly levels, see [Detecting Anomalies and Identifying Advisories in Your Nexus Dashboard](#).

## Advisories

- If you chose **Basic** in the **Mode** area, choose one or more of the following severity levels for advisory statistics that you want to configure for the message bus stream:
  - Critical
  - Major
  - Warning
  - Minor

Or click **Select all** to select all available statistics for the message bus stream.

- If you chose **Advanced** in the **Mode** area:
  - Choose one or more of the following categories for advisory statistics that you want to configure for the message bus stream:
    - Best Practices
    - Field Notices
    - HW end-of-life
    - SW end-of-life
    - PSIRT
  - Choose one or more of the following severity levels for advisory statistics that you want to configure for the message bus stream:
    - Critical
    - Major
    - Warning
    - Minor

Or click **Select all** to select all available categories and statistics for the message bus stream. For more information on advisory levels, see [Detecting Anomalies and Identifying Advisories in Your Nexus Dashboard](#).

## Statistics

There are no differences in the settings in the **Statistics** area when you choose **Basic** or **Advanced** in the **Mode** area.

Choose one or more of the following categories to configure the statistics you want to stream over the message bus.

- Interfaces—streams statistics related to network interfaces, such as traffic volume, error rates, and interface status.

- Protocol—streams protocol-specific statistics, including packet counts, protocol errors, and handshake success rates.
- Resource Allocation—streams data about system resources, such as CPU usage, memory consumption, and bandwidth allocation.
- Environmental—streams environmental metrics such as temperature, humidity, and power supply status.
- Endpoints—streams statistics about connected endpoints, including connection status, data throughput, and session durations.

## Faults

There are no differences in the settings in the **Faults** area when you choose **Basic** or **Advanced** in the **Mode** area.

Choose one or more of the following severity levels for fault statistics that you want to configure for the message bus stream:

- Critical
- Major
- Minor
- Warning
- Info

## Audit Logs

There are no differences in the settings in the **Audit Logs** area when you choose **Basic** or **Advanced** in the **Mode** area.

Choose one or more of the following categories for audit logs that you want to configure for the message bus stream:

- Creation
- Deletion
- Modification

## Syslog

Nexus Dashboard supports the export of anomalies in syslog format. You can use the syslog configuration feature to develop network monitoring and analytics applications on top of Nexus Dashboard, integrate with the syslog server to get alerts, and build customized dashboards and visualizations.

After you choose the fabric where you want to configure the syslog exporter and set up the syslog export configuration, Nexus Dashboard establishes a connection with the syslog server and sends data to the syslog server.

Nexus Dashboard exports anomaly records to the syslog server. With syslog support, you can export

anomalies to your third-party tools even if you do not use Kafka.

## Guidelines and limitations for syslog

If the syslog server is not operational at a certain time, messages generated during that downtime will not be received by a server after the server becomes operational.

### Add syslog server configuration

Follow these steps to add syslog server configuration.

1. Navigate to **Admin > System Settings > General**.
2. In the **Remote streaming servers** area, click **Edit**.

The **Remote streaming servers** page displays.

3. Click **Add server**.

The **Add server** page displays.

4. Choose the **Service** as **Syslog**.
5. Choose the **Protocol**.

You have these options.

- o TCP
- o UDP

6. In the **Name** field, provide the name for the syslog server.
7. In the **Hostname/IP address** field, provide the hostname or IP address of the syslog server.
8. In the **Port** field, specify the port number used by the syslog server.
9. If you want to enable secure communication, check the **TLS** check box.



Before you enable **TLS** you must upload the CA certificate for the syslog destination host to Nexus Dashboard. For more information see, [Upload a CA certificate](#).

### Configure syslog to enable exporting anomalies data to a syslog server

Follow these steps to configure syslog to enable exporting anomalies data to a syslog server.

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

Go to **Manage > Fabrics**.

2. Choose the fabric for which you configure streaming settings.
3. From the **Actions** drop-down list, choose **Edit fabric settings**.

The **Edit fabric-name settings** page displays.

4. In the **Edit fabric-name settings** page, click the **External streaming** tab.
5. Click the **Syslog** tab.

The following details display under **Syslog** tab.

**Edit CS1 Settings** ×

i Some fabric settings can only be modified from the ND Cluster owner standaloneM5PND246

General   Telemetry   **External streaming**

Email   Message bus   **Syslog**

---

**General settings**

**Syslog server\***

Select servers ▼

**Facility**

▼

---

**Collection settings**

**Anomalies**   [Select all](#)   [Clear all](#)

● Critical

● Major

● Warning

[Cancel](#)   [Save](#)

6. Make the necessary configurations in the **General settings** area.
  - a. In the **Syslog server** drop down list, choose a syslog server.

The **Syslog server** drop down list displays the syslog servers that you added in the **System Settings** level. For more information, see [Add syslog server configuration](#).

- b. In the **Facility** field, from the drop-down list, choose the appropriate facility string.

A facility code is used to specify the type of system that is logging the message. For this feature, the **local0-local7** keywords for locally used facility are supported.

7. In the **Collection settings** area, choose the desired severity options.

The options available are **Critical**, **Major**, and **Warning**.

8. Click **Save**.

## Upload a CA certificate

Follow these steps to upload a CA certificate for syslog server **TLS**.

1. Navigate to **Admin > Certificate Management**.
2. In the **Certificate management** page, click the **CA certificates** tab, then click **Add CA certificate**.

You can upload multiple files at a single instance.

3. Browse your local directory and choose the certificate-key pair to upload.

You can upload certificates with the **.pem/.cer/.crt/** file extensions.

4. Click **Save** to upload the selected files to Nexus Dashboard.

A successful upload message appears. The uploaded certificates are listed in the table.

# Additional settings

The following sections provide information for additional settings that might be necessary when editing the settings for a campus VXLAN fabric.

## Layer 3 VNI without VLAN

Following is the upper-level process to enable the Layer 3 VNI without VLAN feature in a fabric:

1. (Optional) When configuring a new fabric, check the **Enable L3VNI w/o VLAN** field to enable the Layer 3 VNI without VLAN feature at the fabric level. The setting at this fabric-level field affects the related field at the VRF level, as described below.
2. When creating or editing a VRF, check the **Enable L3VNI w/o VLAN** field to enable the Layer 3 VNI without VLAN feature at the VRF level. The default setting for this field varies depending on the following factors:
  - o For existing VRFs, the default setting is disabled (the **Enable L3VNI w/o VLAN** box is unchecked).
  - o For newly-created VRFs, the default setting is inherited from the fabric settings, as described above.
  - o 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 is inherited from the fabric setting in the child fabric. You can edit the VRF in the child fabric to change the value, if desired.

See the "Create a VRF" section in [Working with Segmentation and Security for Your Nexus Dashboard VXLAN Fabric](#) for more information.

The VRF attachment (new or edited) then uses the new Layer 3 VNI without VLAN mode if the following conditions are met:

- The **Enable L3VNI w/o VLAN** is enabled at the VRF level
- The switch supports this feature and the switch is running on the correct release (see [Guidelines and limitations for Layer 3 VNI without VLAN](#))

The VLAN is ignored in the VRF attachment when these conditions are met.

## Guidelines and limitations for Layer 3 VNI without VLAN

Following are the guidelines and limitations for the Layer 3 without VLAN feature:

- The Layer 3 VNI without VLAN feature is supported on the -EX, -FX, and -GX versions of the Nexus 9000 switches. When you enable this feature at the VRF level, the feature setting on the VRF will be ignored on switch models that do not support this feature.
- When used in a Campus VXLAN EVPN fabric, this feature is only supported on Cisco Nexus 9000 series switches in that type of fabric. This feature is not supported on Cisco Catalyst 9000 series switches in the Campus VXLAN EVPN fabric; those switches require VLANs for Layer 3 VNI configurations.
- This feature is supported on switches running on NX-OS release 10.3.1 or later. If you enable this feature at the VRF level, the feature setting on the VRF is ignored on switches running an NX-OS

image earlier than 10.3.1.

- When you perform a brownfield import in a Data Center VXLAN EVPN fabric, if one switch configuration is set with the **Enable L3VNI w/o VLAN** configuration at the VRF level, then you should also configure this same setting for the rest of the switches in the same fabric that are associated with this VRF, if the switch models and images support this feature.
- For VXLAN fabric or fabric group deployments with Campus switches, use a VNI range within 4096 to 1,677,721 to align with the best practices for VNI allocation in Campus environments.

## Configuring automatic generation of BGP neighbor description

You can enable auto generation of a BGP neighbor description on BGP sessions between route reflectors and Virtual Tunnel Endpoints (VTEPs), border gateways, and border gateways and a route server. Nexus Dashboard added the **Generate BGP EVPN Neighbor Description** option on the **Edit Fabric Settings** page. Once enabled, Nexus Dashboard generates CLIs for a BGP neighbor description, including the host name of the neighbor. Nexus Dashboard stores the BGP neighbor description in a **bgp\_neighbor\_desc** policy. You can edit the policy with a different description.

You can also configure automatic generation of a BGP neighbor description between border gateways in a VXLAN fabric group.

### Guidelines and limitations for automatic generation of a BGP neighbor description

- In brownfield deployments, the neighbor description on the switch is honored.
- For greenfield deployments, the default for this option is on.
- For fabrics upgraded from a previous release, this option is turned off during an upgrade.
- In a VXLAN fabric group, if you enable the **BGP Neighbor Description on Multi-Site Overlay IFC** option, Nexus Dashboard auto fills the BGP neighbor description in an auto generated multi-fabric overlay inter-fabric link (IFC). You can add a description in the manually created overlay IFC. You can also edit the BGP description field in existing IFCs.
- Auto generation of a BGP neighbor description between a route reflector and a VTEP does not include an eBGP fabric that does not have a route reflector.

### How to configure automatic generation of a BGP neighbor description

Follow these steps to configure automatic generation of a BGP neighbor description.

1. On the **Manage > Fabrics** page, choose the Data Center VXLAN EVPN or the Campus VXLAN EVPN fabric that you want to edit.
2. From the **Actions** drop-down list, choose **Edit fabric settings**.
3. Navigate to the **Fabric Management > Protocols** page in the BGP section.
4. Check the **Generate BGP EVPN Neighbor Description** option.
5. Click **Save**.
6. Perform a **Recalculate and deploy** operation after saving the fabric setting to deploy the configuration to the switches.

## How to configure automatic generation of a BGP neighbor description between border gateways in a VXLAN fabric group

Follow these steps to configure automatic generation of a BGP neighbor description between border gateways or between a border gateway and a route server in a VXLAN fabric group.

1. Create a VXLAN fabric group. For more information, see [Creating Fabrics and Fabric Groups](#).

When creating a VXLAN fabric group, navigate to the **DCI** tab, and check the **BGP Neighbor Description on Multi-Site Overlay IFC** option.

2. On the **Manage > Fabrics** page, choose the fabric that you want to add to the fabric group.
3. From the **Actions** drop-down list, choose **Move Fabric into MSD**.
4. On the **Manage > Fabrics** page, choose the VXLAN fabric group for which you want to edit.

The **Edit *fabric-name* Settings** page displays.

5. Navigate to the **DCI** tab.
6. Ensure that the **Multi-Site Overlay IFC Deployment Method** field displays **directPeering** or **routeServer** as the drop-down option.
7. Add another fabric to the VXLAN fabric group.
8. Click **Save**.
9. Perform a **Recalculate and deploy** operation after saving the fabric settings to deploy the configuration to the switches.

Nexus Dashboard auto generates the inter-fabric link sessions between the border gateways in the VXLAN fabric group that you created.

## Adding Cisco Catalyst 9000 series switches and Nexus 9000 series switches to a Campus VXLAN EVPN fabric

Cisco Catalyst 9000 series switches and Nexus 9000 series switches are discovered using SSH. Before adding the switches to the fabric, it is necessary that you configure the switches for SSH discovery as described in **Before You begin**.

Choose one of the following navigation paths to add switches to the fabric:

- Choose **Manage > Fabrics**. Select the required Campus VXLAN EVPN fabric from the list and choose **Actions > Add Switches**.
- Choose **Manage > Fabrics**. Select the required Campus VXLAN EVPN fabric from the list. Go to the **Switches** tab and then choose **Actions > Add Switches**.
- Choose **Manage > Inventory > Switches** and then choose **Actions > Add Switches**. Click **Choose Fabric**, choose the required Campus VXLAN EVPN fabric, and then click **Select**.

*Before you begin*

- Set the default credentials for the device in the **LAN Credentials Management** window if not already set. To navigate to the **LAN Credentials Management** page from Nexus Dashboard, choose **Admin > Switch Credentials > LAN Credentials Management**.

- For StackWise and StackWise Virtual switches, configure the StackWise/StackWise Virtual settings before adding them to the fabric.
- Run the following SSH commands on the Cisco Catalyst 9000 switch console:

```
switch (config)# hostname <hostname>
switch (config)# ip domain name <domain_name>switch (config)# crypto key generate
rsa
switch (config)# ip ssh time-out 90
switch (config)# ip ssh version 2
switch (config)# line vty 1 4
switch (config-line)# transport input ssh
switch (config)# username admin privilege 15 secret <password>
switch (config)# aaa new-model
switch (config)# aaa authentication login default local
switch (config)# aaa authorization exec default local none
```

Enter values for the following fields:

Field	Description
<b>Seed IP</b>	Enter the IP address of the switch in one of the following formats - "2.2.2.20" or "10.10.10.40-60" or "2.2.2.20, 2.2.2.21".  You can import more than one switch by providing the IP address range. The switches must be properly cabled and reachable from Nexus Dashboard.
<b>Authentication Protocol</b>	Choose the authentication protocol from the drop-down list.
<b>Device Type</b>	Choose <b>IOS XE</b> or <b>NX-OS</b> from the drop-down list. If you select <b>IOS XE</b> , the <b>CAT9K</b> radio button appears which is selected by default.
<b>Username</b>	Enter the username for the switch.
<b>Password</b>	Enter the password for the switch.
<b>Set as individual device write credential</b>	Check the checkbox to set the discovery/read credentials as LAN/write credentials for individual devices.



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

Perform the following steps to add switches to a Campus VXLAN EVPN fabric:

1. Click **Discover Switches**.

The switch details are populated.

Nexus Dashboard supports the import of Cisco Catalyst 9500 switches running in StackWise Virtual. The StackWise Virtual configuration to form a pair of Cisco Catalyst 9500 Switches into a virtual switch has to be in place before the import.

For more information on how to configure StackWise Virtual, see the [Configuring Cisco StackWise Virtual](#) chapter in the *High Availability Configuration Guide (Catalyst 9500 Switches)* for the required release.

2. Check the check boxes next to the switches you want to import.

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



Note that the existing configuration on the switches will be erased after adding the switches to the Campus VXLAN EVPN fabric.

3. Click **Add Switches**.

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

4. (Optional) View the details of the device.

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

*What to do next:*

1. Set the appropriate role. The supported roles are:

Cisco Catalyst 9000 series switches	Nexus 9000 series switches
Leaf	Border gateway
Spine	Border gateway spine
Border	Border gateway super spine

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

After discovering the switches, Nexus Dashboard usually assigns **Leaf** as the default role for Cisco Catalyst 9000 series switches and border gateway for Nexus 9000 series switches.

Optionally, you can form a VPC pair if there are Cisco Nexus 9000 switches with a border gateway role in the fabric.



After setting the switch role, if the switch credentials and the default credentials do not match, the **Discovery Status** or the **Mode** column for the switches displays an error. In such case, set the LAN credentials for the switch. For more information, see the section "LAN Credentials Management" in [Managing Your Device Credentials](#).

2. Recalculate the configurations and deploy the configurations to the switches. Proceed to the next section for the steps to perform recalculate and deploy.

# Recalculating and deploying configurations

To recalculate and deploy the configurations to the switch(es) in the Campus VXLAN EVPN fabric, perform the following steps to recalculate configurations:

*Before you begin:*

Set the role of the switch(es) in the fabric.

1. In Nexus Dashboard, navigate to **Manage > Fabrics**.
2. Click the fabric name to open the **Fabric Overview** page.
3. Choose **Actions > Recalculate and deploy**.

Recalculation of configurations starts on the switch(es).

## Creating VRFs in Campus VXLAN EVPN Fabric

Perform the following steps to create tenant VRFs in a Campus VXLAN EVPN fabric:

1. In Nexus Dashboard, choose **Manage > Fabrics**.
2. From the list of available fabrics, double-click the Campus VXLAN EVPN fabric that you have created in the previous step.

The **Fabric Overview** page appears.

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

The **Create VRF** window appears.

4. Enter the required details in the mandatory fields. Some of the fields are autopopulated with default values. You can make changes, as required.
  - o **VRF Name** - Accept the default value or enter a name for VRF. The VRF name should not contain any white spaces or special characters except underscore (\_), hyphen (-), and colon (:).
  - o **VRF ID** - Accept the default or enter an ID for the VRF.
  - o **VLAN ID** - Specifies the corresponding tenant VLAN ID for the network. Enter an ID for the VLAN. If you want to propose a new VLAN for the network, click **Propose Vlan**.
  - o **VRF Template** - Accept the default autopopulated template or choose another template from the list.

The default template is the template name specified in the **VRF Template** field in the **Advanced** tab on the **Create Fabric** window.

- o **VRF Extension Template** - Accept the default autopopulated template or choose another template from the list.

The default template is the template name specified in the **VRF Extension Template** field in the **Advanced** tab on the **Create Fabric** window.

5. Configure the following fields under **General Parameters**, as needed.

- o **VRF VLAN Name** - Enter the VLAN name for the VRF.
  - o **VRF Description** - Enter a description for the VRF.
  - o **VRF Interface Description**- Enter a description for the VRF interface.
6. Click the **Advanced** tab to optionally specify the advanced profile settings.

Some of the fields in the **Advanced** tab are mentioned here. For more information about configuring the fields in the **Advanced** tab, see "Creating VRF" section in the [About Fabric Overview for LAN Operational Mode Setups](#).

- o **VRF Interface MTU** - Specifies VRF interface MTU.
  - o **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.
  - o **Redistribute Direct Route Map** - Specifies the redistribute direct route map name.
  - o **Max BGP Paths** - Specifies the maximum BGP paths. The valid value range is 1-64 for NX-OS and 1-32 for IOS XE.
  - o **Max iBGP Paths** - Specifies the maximum iBGP paths. The valid value range is 1-64 for NX-OS and 1-32 for IOS XE.
  - o **Advertise Host Routes** - Check the check box to control advertisement of /32 and /128 routes to Edge routers.
  - o **Advertise Default Route** - Check the check box to control advertisement of default route internally.
  - o **Config Static 0/0 Route** - Check the check box to control configuration of static default route.
7. Click **Create** to create the VRF or click **Cancel** to discard the VRF.

A message appears indicating that the VRF is created.

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

### What to do next

1. Attach the VRF.
2. Create a loopback interface and select the **VRF\_LITE** extension.

## Attaching VRFs to switches in Campus VXLAN EVPN fabrics

Follow these steps to attach the VRFs and VRF Lite extensions to the switches in the Campus VXLAN EVPN fabric:

1. On the **VRFs** tab in the **Fabric Overview** window, double-click the VRF that you created in the previous section.

The **VRF Attachments** page opens.

2. Go to the **VRF Attachments** tab and choose the VRF corresponding to the switch by checking the check box next to it.
3. Choose **Actions > Edit**.

The **Extension** page opens.

4. Toggle the knob to **Attach** and click **Save**.

Similarly, you can create a loopback interface, and select the **VRF\_LITE** extension.

For more information about attaching and detaching VRFs, see the section "VRF Attachments" in [About Fabric Overview for LAN Operational Mode Setups](#).

### What to do next

Deploy the configurations as follows:

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

## Creating and deploying networks in Campus VXLAN EVPN fabrics

The next step is to create and deploy networks in Campus VXLAN EVPN fabrics.

### Creating networks for Campus VXLAN EVPN fabrics

Follow these steps to create a network for a Campus VXLAN EVPN fabric:

1. In Nexus Dashboard, choose **Manage > Fabrics**.
2. From the list of available fabrics, click the Campus VXLAN EVPN fabric that you have created in the previous step.

The **Fabric Overview** page appears.

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

The **Create Network** page appears.

4. Enter the required details in the mandatory fields. Some of the fields are autopopulated with default values. You can make changes, as required.

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

Field	Description
<b>Network Name</b>	Specifies the name of the network. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-).

Field	Description
<b>Layer 2 Only</b>	Enables you to create a Layer 2 only network.
<b>VRF Name</b>	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 <b>Create VRF</b> . The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).
<b>VLAN ID</b>	Specifies the corresponding tenant VLAN ID for the network. If you want to propose a new VLAN for the network, click <b>Propose VLAN</b> .
<b>Network Template</b>	Auto-populates the universal template. This is only applicable for leaf switches.
<b>Network Extension Template</b>	Auto-populates the universal extension template. This allows you to extend this network to another fabric. The VRF Lite extension is supported. The template is applicable for border leaf switches.
<b>Generate Multicast IP</b>	If you want to generate a new multicast group address and override the default value, click <b>Generate Multicast IP</b> .

5. Configure the following fields in the **General Parameters** tab:



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

Field	Description
<b>IPv4 Gateway/NetMask</b>	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 a network.  <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>If the same IP address is configured in the <b>IPv4 Gateway</b> and <b>IPv4 Secondary GW1</b> or <b>IPv4 Secondary GW2</b> fields of the network template, Nexus Dashboard 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.</p> </div>
<b>IPv6 Gateway/Prefix List</b>	Specifies one or more IPv6 addresses with subnets.
<b>Vlan Name</b>	Enter a name for the VLAN.
<b>Interface Description</b>	Enter a description for the interface. This interface is a switch virtual interface (SVI).
<b>IPv4 Secondary GW1</b>	Enter the gateway IP address for the additional subnet.

Field	Description
<b>IPv4 Secondary GW2</b>	Enter the gateway IP address for the additional subnet.
<b>IPv4 Secondary GW3</b>	Enter the gateway IP address for the additional subnet.
<b>IPv4 Secondary GW4</b>	Enter the gateway IP address for the additional subnet.

6. Configure the following fields in the **Advanced** tab.

Field	Description
<b>Multicast Group Address</b>	<p>The multicast IP address for the network is auto-populated. Multicast group address is a per fabric instance variable and remains the same for all networks by default. * Configure DHCP relay server fields as follows:</p> <ol style="list-style-type: none"> <li>Under <b>DHCP Relay Server Information</b>, choose <b>Actions &gt; Add</b>.</li> <li>Enter the DHCP relay IP address of the first DHCP server in the <b>Server 1 V4 Address</b> field.</li> <li>Enter the DHCP server VRF ID in the <b>Server VRF</b> field.</li> <li>Click <b>Save</b>.</li> </ol> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>You can configure up to 16 DCHP servers.</p> </div>
<b>Loopback ID for DHCP Relay interface (Min:0, Max:1023)</b>	Enter the loopback ID for DHCP relay interface.
<b>IPv4 TRM Enable</b>	<p>Check the checkbox to enable TRM with IPv4.</p> <p>For more information, see <a href="#">Editing Data Center VXLAN EVPN Fabric Settings</a>.</p>
<b>IPv6 TRM enable</b>	<p>Check the check box to enable TRM with IPv6.</p> <p>For more information, see <a href="#">Editing Data Center VXLAN EVPN Fabric Settings</a>.</p>
<b>L2 VNI Route-Target Both Enable</b>	Check the check box to enable automatic importing and exporting of route targets for all L2 virtual networks. This is applicable only for Cisco Nexus 9000 switches.
<b>Interface Vlan Netflow Monitor</b>	Specifies the Netflow monitor specified for Layer 3 record for the VLAN interface. This is applicable only if Layer 2 Record is not enabled in the Netflow Record for the fabric. This is applicable only for Cisco Nexus 9000 switches.
<b>Vlan Netflow Monitor</b>	Enter the Netflow monitor name defined in the fabric setting for Layer 3 Netflow Record. This is applicable only for Cisco Nexus 9000 switches.
<b>Enable L3 Gateway on Border</b>	Check the check box to enable a Layer 3 gateway on the border switches.

7. Click **Create**.

A message appears indicating that the network is created. The new network appears on the

**Networks** page that comes up.

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

## Attaching networks in a Campus VXLAN EVPN fabric

To attach a network in a Campus VXLAN EVPN fabric from Nexus Dashboard, perform the following steps:

1. In Nexus Dashboard, choose **Manage > Fabrics**.
2. From the list of available fabrics, click the Campus VXLAN EVPN fabric.

The **Fabric Overview** page appears.

3. Navigate to the **Networks** tab and click the network that you created in the previous section to open the **Network Attachments** page.
4. On the **Network attachments** tab, select the required network and choose **Actions > Edit**.

The **Edit Network Attachment** page opens.

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

## Deploying networks in Campus VXLAN EVPN fabrics

To attach networks in a Campus VXLAN EVPN fabric from Nexus Dashboard, perform the following steps:

1. In Nexus Dashboard, choose **Manage > Fabrics**.
2. From the list of available fabrics, double-click the Campus VXLAN EVPN fabric.

The **Fabric Overview** page appears.

3. Navigate to the **Switches** tab, select the switches and choose **Actions > Deploy**.
4. Click **Deploy All** after the configuration preview is complete.
5. Click **Close** after the deployment is complete.

## Creating DCI links for switches in Campus VXLAN EVPN fabrics

You can create a VRF-Lite IFC between a Cisco Catalyst 9000 series switch or a Cisco Nexus 9000 switch with a border role in a Campus VXLAN EVPN fabric, and another switch in a different fabric. The other switch can be a Cisco Nexus 9000 switch in an External fabric, LAN Classic fabric, or Campus VXLAN EVPN fabric. It can also be a Catalyst 9000 switch in an External Fabric or a Campus VXLAN EVPN fabric. The link can be created only from a Campus VXLAN EVPN fabric. The other switch can be an External, Classic LAN, Campus VXLAN EVPN, or a Data Center VXLAN EVPN fabric. The link can be created only from the Campus VXLAN EVPN fabric.

For more information, see the section "Links" in [About Fabric Overview for LAN Operational Mode](#)



When creating DCI links for a Campus VXLAN EVPN fabric, auto-deploy is supported only if the destination device is a Cisco Nexus 9000 switch.

Follow these steps for creating links for a Campus VXLAN EVPN fabric:

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

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

2. Choose **Actions > Create**.

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

3. From the **Link Type** drop-down box, choose **Inter-Fabric**.
4. From the **Link Sub-Type** drop-down list, choose **VRF\_LITE**.
5. In the **Link Template** field, ensure **ext\_fabric\_setup** template is auto populated for VRF\_LITE IFC.

The templates are autopopulated with corresponding pre-packaged default templates that are based on your selection. The template to use for **VRF\_LITE** IFC is **ext\_fabric\_setup**.

6. From the **Source Fabric** drop-down list, choose the campus VXLAN EVPN fabric.
7. From the **Destination Fabric** drop-down list, choose a destination fabric.
8. Choose the **Source Device** and Ethernet interface that connects to the destination device.
9. Choose the **Destination Device** and Ethernet interface that connects to the source device.
10. Enter values in other fields, as required. For more information about configuring the fields in the **Default VRF** tab, see the section "Establishing inter-fabric connectivity using VRF Lite" in [Working with Connectivity in Your Nexus Dashboard LAN Fabrics](#).
11. Click **Save**.

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

Choose **VRF\_LITE** as the link sub-type and **ext\_fabric\_setup** template for VRF\_LITE IFC. To complete the procedure, repeat step 4 to step 10 mentioned above.

# 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)

© 2017–2026 Cisco Systems, Inc. All rights reserved.

## **Americas Headquarters**

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

<https://www.cisco.com>

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883