

Editing External Fabric Settings, Release 4.3.1

Table of Contents

New and changed information . . . 1
Editing external fabric settings . . . 2
General . . . 3
Fabric Management . . . 5
  General Parameters . . . 5
  Advanced . . . 6
  Resources . . . 10
  Configuration Backup . . . 10
  Bootstrap . . . 11
  Flow Monitor . . . 13
Telemetry . . . 16
  Edit configuration settings . . . 17
  Edit NAS settings . . . 17
  Disable Telemetry . . . 18
  Perform force disable telemetry on your fabric . . . 18
  NAS . . . 18
  Guidelines and limitations for network attached storage . . . 19
  Add network attached storage to export flow records . . . 19
    Add NAS to Nexus Dashboard . . . 19
    Add the onboarded NAS to Nexus Dashboard . . . 20
  Flow collection . . . 22
  Understanding flow telemetry . . . 22
    Flow telemetry guidelines and limitations . . . 22
    Configure flows . . . 24
    Monitor the subnet for flow telemetry . . . 27
  Understanding Netflow . . . 28
    Understanding Netflow types . . . 28
    Netflow guidelines and limitations . . . 29
    Configure Netflow . . . 30
  Understanding sFlow . . . 31
    Guidelines and limitations for sFlow . . . 31
    Configure sFlow telemetry . . . 32
External streaming . . . 33
  Configure external streaming settings . . . 33
  Guidelines and limitations . . . 34
    Guidelines and limitations in NX-OS fabrics . . . 34
  Email . . . 34
  Message bus . . . 36
    Add Kafka broker configuration . . . 36
    Configure Kafka exports in fabric settings . . . 38
    Anomalies . . . 40

Advisories . . . 41
Statistics . . . 41
Faults . . . 42
Audit Logs . . . 42
Syslog . . . 42
Guidelines and limitations for syslog . . . 43
Add syslog server configuration . . . 43
Configure syslog to enable exporting anomalies data to a syslog server . . . 43
Additional settings . . . 45
External Fabrics . . . 45
Guidelines and Limitations . . . 45
Move an External Fabric Under an MSD Fabric . . . 47
External Fabric Depiction in an MSD Fabric Topology . . . 48
Add switches to the External fabric . . . 48
Switch Settings for External Fabrics . . . 49
Discover new switches . . . 50
Adding Non-Nexus Devices to External Fabrics . . . 53
Configuration Compliance in External Fabrics . . . 54
Special Configuration CLIs Ignored for Configuration Compliance . . . 55
Managing Cisco IOS-XR Devices using NDFC . . . 55
Configuring Non-Nexus Devices for Discovery . . . 56
Discovering Non-Nexus Devices in an External Fabric . . . 60
Managing Non-Nexus Devices to External Fabrics . . . 63
Understanding switches running SONiC . . . 63
Switches running SONiC: On AI Routed and Routed fabrics . . . 63
Switches running SONiC: External fabrics . . . 65
Switches running SONiC: General information . . . 65
Troubleshooting issues . . . 66
Creating a vPC Setup . . . 66
Undeploying a vPC Setup . . . 67
Precision Time Protocol for External Fabrics . . . 68
Copyright . . . 71

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 4.3.1	Support for switches running SONiC images	Support is now available for certain switches running SONiC (Software for Open Networking in Cloud) images in these Nexus Dashboard fabric types: - AI Routed - Routed - External (this article) For more information, see Understanding switches running SONiC.

Nexus Dashboard 4.3.1	Support for new Cisco Catalyst and NX-OS switches	Beginning with Nexus Dashboard 4.3.1, support is added for the following Catalyst and NX-OS switches in External fabric:
		- Cisco 8400 Series Secure Routers:
		  - Cisco Secure Router 8475-G2
		  - Cisco Secure Router 8455-G2
		- Cisco Catalyst 9000 Series Switches:
		  - C9300L-24P-4G
		  - C9300L-48-4G
		  - C9200CX-12P-2X2G-E
		- Cisco N9164E-NS4-O switch running NX-OS release 10.6.2(n) or later For more information, see Guidelines and Limitations.

Nexus Dashboard 4.3.1	Image management in monitor mode	Beginning with Nexus Dashboard 4.3.1, you can perform image management operations for supported fabrics while the fabrics are in monitor mode. For more information, see General Parameters.
Nexus Dashboard 4.3.1	Cisco Nexus 9000 Series name change	Cisco Nexus 9000 Series is now the Cisco N9000 Series.

Editing external fabric settings

An **external fabric** is a type of fabric that is used to automate provisioning of a network that might include Cisco NX-OS, IOS-XE, IOS-XR, SONiC, or third-party devices for monitoring or provisioning. This includes use cases for external connectivity and multi-site interconnectivity (IPNs or ISNs).

When you first create an external 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 external fabric.

1. Navigate to the main Fabrics window:

   Manage > Fabrics

2. Locate the external fabric that you want to edit.

   External fabrics are shown with **External and Inter-Fabric Connectivity** in the **Type** column.

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

   The **Edit *fabric_name* Settings** window appears.

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

General

Use the information in this section to edit the settings in the **General** window for your external fabric.

Change the general parameters that you configured previously for the external fabric, if necessary, or click another tab to leave these settings unchanged.

Fabric type	Description
Name	The name for the fabric. This field is not editable.
Type	The fabric type for this fabric. This field is not editable.
Location	Change the location for the fabric, if necessary.
BGP ASN	Change the BGP autonomous system number (ASN) for the fabric, if necessary.  When you update the **BGP ASN** field in a fabric configured with External Connectivity, the change is not immediately reflected in the deployment preview nor applied during the **Recalculate & Deploy** operation. The updated ASN takes effect only after the affected switch is removed from the fabric and re-onboarded.
License tier for fabric	Change the licensing tier for the fabric, if necessary: · **Essentials** · **Advantage** · **Premier** Click on the information icon (?) next to License tier to see what functionality is enabled for each license tier.
Fabric monitor mode	Check this check box to monitor the fabric only. Nexus Dashboard Fabric Controller does not deploy any configuration to a fabric in monitor mode.
Enable fabric software update	Check this check box to enable image management for switch updates on monitored fabrics. This option is available only when **Fabric monitor mode** is enabled. This option enables image management for switch updates even for monitored mode fabrics.
Enabled features	Check the box to enable **Telemetry** for the fabric. This is the equivalent of enabling the Nexus Dashboard Insights service in previous releases.
Telemetry collection	This option becomes available if you choose to enable **Telemetry** in the **Enabled features** field above. Choose either **Out-of-band** or **In-band** for telemetry collection.
Telemetry streaming via	This option becomes available if you choose to enable **Telemetry** in the **Enabled features** field above. Choose either **IPv4** or **IPv6** for telemetry streaming.

Fabric type	Description
Telemetry VRF	This option becomes available if: - You enabled **Telemetry** in the **Enabled features** field, and - You clicked **Advanced** in the **Configuration Mode** field. Enter the appropriate VRF instance in this field.
Telemetry source interface	This option becomes available if you enabled **Telemetry** in the **Enabled features** field,. Enter the source interface for telemetry streaming.
Security domain	Change the security domain for the fabric, if necessary.

Fabric Management

Use the information in this section to edit the settings in the **Fabric management** page for your external fabric. The following sections describe each tab and its respective fields. These tabs include the fabric-level parameters.

- General Parameters
- Advanced
- Resources
- Configuration Backup
- Bootstrap
- Flow Monitor

General Parameters

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

Field	Description
Fabric monitor mode	Clear the check box if you want Nexus Dashboard Fabric Controller to manage the fabric. Keep the check box checked to enable a monitor-only external fabric. From the Cisco Nexus Dashboard Fabric Controller Release 12.1.1e, you can also add Cisco 8000 Series Routers to external fabrics both in managed mode and monitored mode. When you create an Inter-Fabric Connection (IFC) from a VXLAN fabric to this external fabric, the BGP AS number is referenced as the external or neighbor fabric AS number. When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches. If you click **Deploy Config**, it displays an error message. The configurations must be pushed for non-Nexus devices before you discover them in the fabric. You cannot push configurations in the monitor mode.

Field	Description
Enable fabric software update	Check this check box to allow switch upgrade-related operations for supported external fabrics in monitor mode. This option is available only when **Fabric monitor mode** is enabled. The check box is not enabled by default. When this option is enabled, Nexus Dashboard can perform image management operations and supported switch actions, such as reloading switches, running Show commands, changing switch mode, and updating fabric agent packages for Connectivity Analysis. Other fabric edit and deployment operations remain unavailable while the fabric is in monitor mode.
Enable Performance Monitoring (For NX-OS and IOS XE Switches Only)	Check this check box to enable performance monitoring on NX-OS switches only. Ensure that you do not clear interface counters from the command-line interface of the switches. Clearing interface counters can cause the **Performance Monitor** to display incorrect data for traffic utilization. If you must clear the counters and the switch has both clear counters and clear counters snmp commands (not all switches have the clear counters snmp command), ensure that you run both the main and the SNMP commands simultaneously. For example, you must run the clear counters interface ethernet slot/port command followed by the clear counters interface ethernet slot/port snmp command. This can lead to a one time spike.

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

Advanced

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

Field	Description
Power Supply Mode	Choose the appropriate power supply mode.
Enable MPLS Handoff	Check this check box to enable the MPLS Handoff feature. For more information, see Overview of VXLAN EVPN to SR-MPLS and MPLS LDP interconnection
Underlay MPLS Loopback Id	Specifies the underlay MPLS loopback ID. The default value is 101.
Enable AAA IP Authorization	Enables AAA IP authorization, after IP authorization is enabled on the AAA server.
Enable NDFC as Trap Host	Check this check box to enable the Nexus Dashboard Fabric Controller as a trap host.
Enable CDP for Bootstrapped Switch	Check the check box to enable CDP for the bootstrapped switch.

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

Field	Description
Inband Mgmt	For External and Classic LAN fabrics, this knob enables Nexus Dashboard Fabric Controller to import and manage switches with inband connectivity (reachable over switch loopback, or a routed interface, or SVI interfaces), in addition to management of switches with out-of-band connectivity (reachable over the switch mgmt0 interface). The only requirement is that for inband-managed switches, there should be IP reachability from Nexus Dashboard Fabric Controller to the switches over the Nexus Dashboard data interface, also known as an inband interface. For this purpose, static routes may be needed on the Nexus Dashboard Fabric Controller, that in turn can be configured from **Admin > System Settings > Routes**. After enabling inband management, during discovery, provide the IPs of all the switches to be imported using inband management and set maximum hops to 0. Nexus Dashboard Fabric Controller has a precheck that validates that the inband-managed switch IPs are reachable over the Nexus Dashboard data interface. After completing the precheck, Nexus Dashboard Fabric Controller discovers and learns about the interface on that switch that has the specified discovery IP in addition to the VRF that the interface belongs to. As part of the process of switch import/discovery, this information is captured in the baseline intent that is populated on the Nexus Dashboard Fabric Controller. For more information, see the section "Inband management in External fabrics and LAN Classic fabrics" in Configuring Inband Management and Out-of-Band PnP.  Bootstrap or POAP is only supported for switches that are reachable over out-of-band connectivity, that is, over switch mgmt0. The various POAP services on the Nexus Dashboard Fabric Controller are typically bound to the eth1 or out-of-band interface. In scenarios, where Nexus Dashboard Fabric Controller eth0/eth1 interfaces reside in the same IP subnet, the POAP services are bound to both interfaces.
Enable Precision Time Protocol (PTP)	Enables PTP across a fabric. When you check this check box, PTP is enabled globally and on core-facing interfaces. You can also edit the **PTP Source Loopback Id** and **PTP Domain Id** fields. For more information, see the section Precision Time Protocol for External Fabrics.
PTP Source Loopback Id	Specifies the loopback interface ID loopback that is used as the source IP address for all Precision Time Protocol (PTP) packets. The valid values range from 0-1023. The PTP loopback ID cannot be the same as RP, Phantom RP, NVE, or the MPLS loopback ID. Otherwise, an error is generated. The PTP loopback ID can be the same as the BGP loopback or user-defined loopback that is created from Nexus Dashboard Fabric Controller. If the PTP loopback ID is not found during a **Save & Deploy**, the following error is generated: Loopback interface to use for the PTP source IP is not found. Please create a PTP loopback interface on all the devices to enable the PTP feature.

Field	Description
PTP Domain Id	Specifies the PTP domain ID on a single network. The valid values range from 0-127.
Enable Real Time Interface Statistics Collection	Valid for NX-OS only. Check the box to enable the collection of real time interface statistics.
Interface Statistics Load Interval	Enter the time, in seconds, for the interface statistics load interval (Min:5, Max:300).
CoPP Profile	Choose the appropriate control plane policing (CoPP) profile for the fabric. These profiles are available. • dense • lenient • moderate • strict • manual The **manual** option is chosen by default.  In general, a fabric-wide CoPP policy is applied to Nexus switches. If **manual** option is chosen, a customized CoPP profile policy must be defined separately.
 Freeform is not supported for switches running on the SONiC image. The config replay is instead provided to you to upload the JSON format configurations to merge or replace the running configurations on SONiC switches.	
Fabric Freeform	You can apply configurations globally across all the devices that are discovered in the external fabric using this freeform field. The devices in the fabric should belong to the same device type and the fabric should not be in monitor mode. The different device types are: • NX-OS • IOS-XE • IOS-XR • Others Depending on the device type, enter the configurations accordingly. If some of the devices in the fabric do not support these global configurations, they go out-of-sync or fail during the deployment. Hence, ensure that the configurations you apply are supported on all the devices in the fabric or remove the devices that do not support these configurations.
AAA Freeform Config	You can apply AAA configurations globally across all devices that are discovered in the external fabric using this freeform field.

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

Resources

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

Field	Description
Subinterface Dot1q Range	The subinterface 802.1Q range and the underlay routing loopback IP address range are autopopulated.
Underlay MPLS Loopback IP Range	Specifies the underlay MPLS SR or LDP loopback IP address range. The IP range should be unique, that is, it should not overlap with IP ranges of the other fabrics.

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

Configuration Backup

The fields in the **Configuration Backup** tab are described in the following table.

Field	Description
Hourly Fabric Backup	Check the check box to enable an hourly backup of the fabric configurations and the intent. You can enable an hourly backup for fresh fabric configurations and the intent as well. If there is a configuration push in the previous hour, Nexus Dashboard Fabric Controller takes a backup. In case of the external fabric, the entire configuration on the switch is not converted to intent on Nexus Dashboard Fabric Controller as compared to the VXLAN fabric. Therefore, for the external fabric, both the intent and the running configuration are backed up. **Intent** refers to configurations that are saved in Nexus Dashboard Fabric Controller, but yet to be provisioned on the switches. The hourly backups are triggered during the first 10 minutes of the hour.
Scheduled Fabric Backup	Check the check box to enable a daily backup. This backup tracks changes in running configurations on the fabric devices that are not tracked by configuration compliance.

Field	Description
Scheduled Time	Specify the scheduled backup time in a 24-hour format. This field is enabled if you check the **Scheduled Fabric Backup** check box. Check both the check boxes to enable both backup processes. The backup process is initiated after you click **Save**. The scheduled backups are triggered exactly at the time that you specify with a delay of up to two minutes. The scheduled backups are triggered regardless of the configuration deployment status. You can also initiate the fabric backup on the **Overview > Topology** page. Click **Backup Fabric** in the **Actions** drop-down list. The backups contain the running configuration and the intent that is pushed by Nexus Dashboard Fabric Controller. Configuration compliance forces the running configuration to be the same as the Nexus Dashboard Fabric Controller configuration. Note that for the external fabric, only some configurations are part of the intent and the remaining configurations are not tracked by Nexus Dashboard Fabric Controller. Therefore, as part of the backup, both Nexus Dashboard Fabric Controller intent and the running configuration from the switch are captured.



When an individual fabric is a member of a Multi-Site Domain (MSD) fabric, the fabric backup option for that individual fabric is disabled. Perform fabric backups for MSD member fabrics at the MSD level. This ensures a comprehensive backup of the entire multi-site deployment.

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

Bootstrap

You can configure automatic Cisco Plug n Play (PnP) IP assignment of Cisco Catalyst 9000 Series switches in an External Connectivity Network or a Custom Network fabric by configuring the following.

1. Navigate to the **Admin > Certificate Management > Bootstrap Certificates** tab, and click the **Upload Certificate** button.

   The **Upload Certificate - Bootstrap Server** dialog displays.

2. Drag and drop your bootstrap certificate file to the dialog box or browse to the location of your file.

   The following are the accepted file types: .pem, .cer, .key, or .crt.

3. Enter your password for your bootstrap server and click **Upload**.
4. Ensure that you configure the **Bootstrap Script Download Protocol** field in **Admin > System Settings > LAN-Fabric** as **http**.

5. On the **Create Fabric > Bootstrap** or the **Edit Fabric > Bootstrap** tab, configure the following:
    a. Check the **Enable Bootstrap (For NX-OS and IOS-XE(Cat9K) Switches Only)** check box.
    b. Check the **Enable Local DHCP Server** check box.
    c. Check the **Enable Plug n Play for Cat9K** check box.

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

Field	Description
Enable Bootstrap (For NX-OS and IOS-XE(Cat9K) Switches Only	Check this check box to enable the bootstrap feature for NX-OS and IOS-XE Cisco Catalyst 9000 Series switches. After you enable bootstrap, you can enable the DHCP server for automatic IP address assignment.
Enable Inband POAP	Check this check box to enable inband POAP.  You must enable **Inband Mgmt** on the **Advanced** tab to enable this option.
Enable Local DHCP Server	Enable the **Local DHCP Server** check box and enter details for the remaining mandatory fields. From Cisco NDFC Release 12.1.1e, you can choose inband POAP or out-of-band POAP for external fabrics.
Enable Plug n Play for Cat9K	Check this check box to enable PnP for automatic IP assignment for Cisco Catalyst 9000 Series switches.  Cisco PnP is supported for Cisco Catalyst 9000 Series switches only.
DHCP Version	Choose DHCPv4 or DHCPv6 from this drop-down list. When you choose DHCPv4, **Switch Mgmt IPv6 Subnet Prefix** field is disabled. If you choose DHCPv6, the **Switch Mgmt IP Subnet Prefix** is disabled.  Cisco Nexus Dashboard Fabric Controller IPv6 POAP is not supported with Cisco Nexus 7000 Series switches. Cisco N9000 and Cisco Nexus 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 Layer 3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other than /64 are not supported.  NDFC supports Cisco PnP for IPv4 only. There is no support for IPv6. If you do not check this check box, Nexus Dashboard Fabric Controller uses the remote or external DHCP server for automatic IP address assignment.
Domain name	Specify the domain name for the DHCP server PnP block.

Field	Description
DHCP Scope Start Address and **DHCP Scope End Address**	Specifies the first and last IP addresses of the IP address range to be used for the switch out-of-band POAP.
Switch Mgmt Default Gateway	Specifies the default gateway for the management VRF on the switch.
Switch Mgmt IP Subnet Prefix	Specifies the prefix for the mgmt0 interface on the switch. The prefix range is 8-30.
Switch Mgmt IPv6 Subnet Prefix	Specifies the IPv6 prefix for the Mgmt0 interface on the switch. The prefix should be from 112 through 126. This field is editable if you enable IPv6 for DHCP.
Enable AAA Config	Check this check box to include AAA configs from Advanced tab during device bootup.
Bootstrap Freeform Config (Optional)	Enter other commands as needed. For example, if you are using AAA or remote authentication-related configurations, add these configurations in this field to save the intent. After the devices boot up, they contain the intent that is defined in the **Bootstrap Freeform Config** field. Copy-paste the running-config to a **freeform config** field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. For more information, see the section "Enable freeform configurations on fabric switches" in Working with Inventory in Your Nexus Dashboard LAN or IPFM Fabrics.
DHCPv4 Multi Subnet Scope	Specifies the field to enter one subnet scope per line. This field is editable after you check the **Enable Local DHCP Server** check box. The format of the scope should be defined as: **DHCP Scope Start Address, DHCP Scope End Address, Switch Management Default Gateway, Switch Management Subnet Prefix** Example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

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

Flow Monitor

The fields in the **Flow Monitor** tab are described in the following table.

Field	Description
Enable NetFlow	Check this check box to enable NetFlow on VTEPs for this fabric. By default, NetFlow is disabled. On Enable, NetFlow configuration will be applied to all VTEPS that support NetFlow.  When NetFlow is enabled on the fabric, you can choose not to have NetFlow on a particular switch by having a dummy no_netflow PTI. If NetFlow is not enabled at the fabric level, an error message is generated when you enable NetFlow at the interface, network, or VRF level. For information about NetFlow support for Cisco NDFC, see the "Configuring Netflow support" section in Creating Fabrics and Fabric Groups.
NetFlow Exporter	Click **Actions > Add** to add one or more NetFlow exporters. This exporter is the receiver of the NetFlow data. The fields on this area are: • **Exporter Name** – Specifies the name of the exporter. • **IP** – Specifies the IP address of the exporter. • **VRF** – Specifies the VRF over which the exporter is routed. • **Source Interface** – Enter the source interface name. • **UDP Port** – Specifies the UDP port over which the NetFlow data is exported. Click **Save** to configure the exporter. Click **Cancel** to discard. You can also choose an existing Netflow exporter and choose **Actions > Edit** or **Actions > Delete** to perform the relevant actions.

NetFlow Record	Click **Actions > Add** to add one or more NetFlow records. The fields on this area are:
	
	• **Record Name** – Specifies the name of the record.
	• **Record Template** – Specifies the template for the record. Enter one of the record templates names. You can create custom NetFlow record templates. Custom record templates that are saved in the template library are available for use here.
	  ○ **netflow_ipv4_record** – to use the IPv4 record template.
	  ○ **netflow_l2_record** – to use the Layer 2 record template.
	• **Is Layer 2 Record** – Check this check box if the record is for a Layer 2 NetFlow.
	
	Click **Save** to configure the report. Click **Cancel** to discard. You can also choose an existing record and choose **Actions > Edit** or **Actions > Delete** to perform the relevant actions.

Field	Description
NetFlow Monitor	Click **Actions > Add** to add one or more NetFlow monitors. The fields on this area are: · **Monitor Name** – Specifies the name of the monitor. · **Record Name** – Specifies the name of the record for the monitor. · **Exporter1 Name** – Specifies the name of the exporter for the NetFlow monitor. · **Exporter2 Name** – (optional) Specifies the name of the secondary exporter for the NetFlow monitor. The record name and exporters referred to in each NetFlow monitor must be defined in the **Netflow Record** and **Netflow Exporter** areas. Click **Save** to configure the monitor. Click **Cancel** to discard. You can also choose an existing Netflow monitor and choose **Actions > Edit** or **Actions > Delete** to perform the relevant actions.
NetFlow Sampler	Click **Actions > Add** to add one or more NetFlow samplers. The fields on this area are: · **Sampler Name** – Specifies the name of the Netflow sampler. · **Number of Samples** – Specifies the number of Netflow samples. · **Number of Packets in Each Sampling** – Specifies the number of packets in each Netflow sampling.  The **Netflow Sampler** is applicable to Cisco Nexus 7000 Series switches only. Click **Save** to configure the Netflow sampler. Click **Cancel** to discard. You can also choose an existing Netflow sampler and choose **Actions > Edit** or **Actions > Delete** to perform the relevant actions.

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

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** page.
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.
   - 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.
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.



Beginning with Nexus Dashboard 4.3.1, flows are exported to a path with the unified Nexus Dashboard version in the directory name instead of the NIR version.

```
└── NDI-<VERSION>-FLOW-JSON/
    └── fabricName=<fabricName>/
        └── year=2022/
            ├── month=01/
            │   └── date=01/
            │       ├── hour=01/
            │       │   ├── 52170795-0b94-481c-800a-c47f0fa41fac.json
            │       │   └── fa92c70c-96fc-4e32-ac76-324bdd5139d4.json
            │       └── hour=23/
            │           ├── 737f4292-bf29-4630-bdd9-ccb80885ddc1.json
            │           └── 68b434d9-0957-4fe4-be01-e0688cb4336d.json
            └── month=02/
                └── date=20/
                    ├── hour=10/
                    │   ├── e05ce8fb-88af-45db-8c52-4b00e1841b16.json
                    │   └── 6fd2b652-dfe1-430e-905a-020abd399e3e.json
                    └── hour=23/
                        ├── eeb6784a-33a0-4ae3-b13e-db4db93fe48b.json
                        └── b289c75e-a709-4284-a018-b38ab101d90f.json
```

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" :" myapic" ," terminalTs" :1688537547433," originTs" :1688537530376," srcIp" :" 2000:201:1:1::1" ," dstIp" :" 2000:201:1:1::3" ," srcPort" :1231," dstPort" :1232," ingressVrf" :" vrf1" ," egressVrf" :" vrf1" ," ingressTenant" :" FSV1" ," egressTenant" :" FSV1" ," protocol" :" UDP" }

{" fabricName" :" myapic" ," terminalTs" :1688537547378," originTs" :1688537530377," srcIp" :" 201.1.1.127" ," dstIp" :" 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" :" myapic" ," terminalTs" :1688538023562," originTs" :1688538010527," srcIp" :" 201.1.1.121" ," dstIp" :" 201.1.1.127" ," srcPort" :0," dstPort" :0," ingressVrf" :" vrf1" ," egressVrf" :" 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," ingressBurst" :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," events" :" [[\\\" 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.
- You can enable Flow Telemetry at the fabric level in Nexus Dashboard, even if only a subset of switches in the fabric supports the feature. Enabling Flow Telemetry at the fabric level does not

require all switches to support it; provided at least one switch supports Flow Telemetry, you can enable the feature for the entire fabric.

- 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.
- The **Fabric Overview** page displays the correct PTP status only when **Flow telemetry** or **Traffic analytics** is enabled under **Flow collection**. Enabling PTP in the fabric settings alone does not enable PTP monitoring.
- 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.
- For inter-fabric flows that traverse an external fabric and return to the source fabric (a U-turn path), Nexus Dashboard displays the egress interface as N/A when the external fabric contains Cisco N9K-C9364C switches. Nexus Dashboard cannot capture egress port metadata for these flows. The remote port value remains valid.
- 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.
- 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.
  - Re-enable Flow Telemetry.
  - 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** page 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** page 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** page.
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.
   - VRF
   - Physical interface
   - Port channel

VRF

Follow these steps 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

Follow these steps 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

Follow these steps 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.

Cisco N9000 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 N9000 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 Cisco Nexus 7000 and Cisco N9000 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 N9000 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
L2 switchport Interface (Ethernet/Port-channel)	Yes	Yes	Yes	Yes	No	No

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 N9000 Series fabric modules:
  - Cisco N9508-R fabric module (N9K-C9508-FM-R)
  - Cisco N9504-R fabric module (N9K-C9504-FM-R)
- sFlow utilizes statistical sampling to monitor traffic. As a result, ingress and egress data are reported as separate records and may not be available for every packet flow.
- To configure sFlow on fabric switches, see the **Configuring sFlow** section in the Cisco N9000

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:
    - Choose **sFlow**.
    - 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.

   - Email
   - Message bus
   - 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.

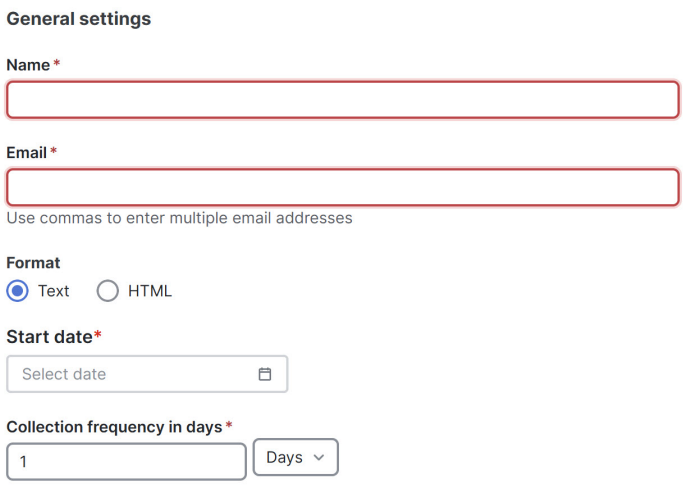
4. In the **Edit *fabric-name* settings** page, click the **External streaming** tab.
5. Click the **Email** tab.
6. Review the information provided in the **Email** tab for already-configured email configurations.

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

Field	Description
Name	The name of the email configuration.
Email	The email addresses used in the email configuration.
Start time	The start date used in the email configuration.
Frequency	The frequency in days or weeks set in the email configuration.
Anomalies	The severity level for anomalies and advisories set in the email configuration.
Advisories	
Risk and conformance reports	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.



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



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
Certificate name	The name of the CA certificate.
Certificate details	The details of the CA certificate.
Attached to	The CA signed certificate attached to Nexus Dashboard.
Expires on	The Expiry date and time of the CA certificate.
Last updated time	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
Certificate name	The name of the Client certificate.
Certificate details	The details of the Client certificate.
Attached to	The feature to which the system certificate is attached to, in this case, the message bus.
Expires on	The Expiry date and time of the CA certificate.
Last updated time	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
Message bus stream	The name of the message bus stream configuration.
Collection type	The collection type used by the message bus stream.
Mode	The mode used by the message bus stream.
Anomalies	The severity level for anomalies and advisories set in the message bus stream configuration.
Advisories	
Statistics	The statistics that were configured for the message bus stream.
Faults	The severity level for faults set in the message bus stream configuration.
Audit Logs	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.

- **Alerts and events**: This is the default setting. Continue to Step 7, if you choose **Alerts and events**.
- **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 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
        - Security advisory
    - 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 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.

   - TCP
   - 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

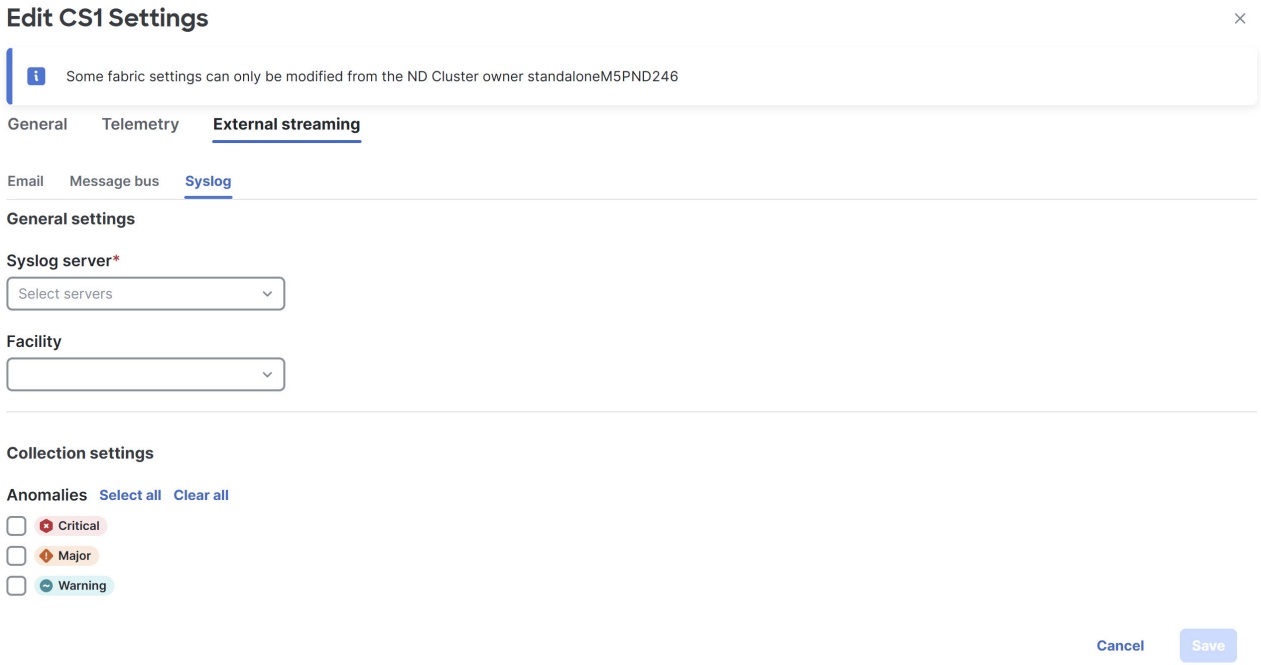
Follow these steps to configure syslog to enable exporting anomalies data to a syslog server.

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.



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 an external fabric.

External Fabrics

You can add switches to an external fabric.

Guidelines and Limitations

- NDFC will not generate "no router bgp". If you want to change it, go to the switch and do a "no feature bgp" followed by a re-sync, if you don't have anything and want to update the ASN.
- The external fabric is a monitor-only or managed-mode fabric.
- Beginning with Cisco Nexus Dashboard Fabric Controller release 12.2.2, you can enable Cisco PnP for automatic IP assignment for Cisco Catalyst 9000 Series switches for an External Connectivity Network or a Custom Network fabric. For more information, see Bootstrap.
- From Cisco Nexus Dashboard Fabric Controller release 12.0.1, Cisco IOS-XR family devices Cisco ASR 9000 Series Aggregation Services routers and Cisco Network Convergence System (NCS) 5500 Series switches are supported in an external fabric in managed mode and monitor mode. NDFC generates and pushes configurations to these switches, and configuration compliance is also enabled for these switches.
- From Cisco Nexus Dashboard Fabric Controller release 12.1.1e, you can also add Cisco 8000 Series routers to external fabrics both in managed mode and monitored mode, and configuration compliance is also supported.
- You can import, remove, and delete switches for an external fabric.
- For Inter-Fabric Connection (IFC) cases, you can choose Cisco 9000, 7000 and 5600 Series switches as destination switches in an external fabric.
- You can use non-existing switches as destination switches.
- The template that supports an external fabric is the External Connectivity Network fabric.
- If an external fabric is an MSD fabric member, then the MSD topology screen displays the external fabric with its devices, along with the member fabrics and their devices.

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

- Nexus Dashboard retains configured underlay IP addresses on switches after you remove an external fabric from a fabric group or MSD. To prevent overlapping network errors, such as when re-adding a fabric to a group or when the system does not detect ISN-to-ACI links, perform a **Recalculate & Deploy** on the external fabric immediately after removal to clear stale IP configurations. You must manually remove the IPs from the interfaces if conflicts persist.
- You can set up a Multi-Site or a VRF-Lite IFC by manually configuring the links for the border devices in the VXLAN fabric or by using an automatic Deploy Border Gateway Method or VRF Lite IFC Deploy Method. If you are configuring the links manually for the border devices, we recommend using the Core Router role to set up a Multi-Site eBGP underlay from a Border Gateway device to a Core Router and the Edge Router role to set up a VRF-Lite IFC from a border

device to an edge device.

- If you are using the Cisco Nexus 7000 Series switch with Cisco NX-OS release 6.2(24a) on the LAN Classic or an External Connectivity Network fabric, make sure to enable **AAA IP Authorization** in the fabric settings.
- In NDFC release 12.2.1, Cisco N9800 switches are supported only with roles of Spine and Super-Spine in a VXLAN EVPN fabric with VXLAN OAM disabled. If a Cisco N9800 switch is positioned as a route server (if it has been assigned with a Core Router role) in an External Connectivity Network fabric, that Cisco N9800 switch must be running on NX-OS release 10.4.2 or later.
- You can discover the following non-Nexus devices in an External Connectivity Network fabric:
  - IOS-XE family devices: Cisco CSR 1000v, Cisco ASR 1000 Series routers, and Cisco Catalyst 9000 Series switches
  - IOS-XR family devices: ASR 9000 Series routers, IOS XR release 6.5.2, and Cisco NCS 5500 Series routers, IOS XR release 6.5.3
  - Arista 4.2 (any model)
- Configure all the non-Nexus devices, except Cisco CSR 1000v, before adding them to an external fabric.
- You can configure non-Nexus devices as borders. You can create an IFC between a non-Nexus device in an external fabric and a Cisco Nexus device in an easy fabric. The interfaces supported for these devices are:
  - Routed
  - Subinterface
  - Loopback
- You can configure a Cisco ASR 1000 Series router and a Cisco Catalyst 9000 Series switch as an edge router, set up a VRF-Lite IFC, and connect it as a border device with an easy fabric.
- Before a VDC reload, discover Admin VDC in the fabric. Otherwise, the reload operation does not occur.
- You can connect a Cisco data center to a public cloud using Cisco CSR 1000v. For more information, see the "Setting Up the Infra Configuration for Hybrid Cloud and Multi-Cloud Connectivity Deployment" section in the Hybrid Cloud Connectivity Deployment for Cisco NX-OS Guide for the use case.
- In an external fabric, when you add the **switch_user** policy and provide the username and password, the password must be an encrypted string that is displayed in the show run command.

  For example:

  ```
  username admin password 5
  $5$I4sapkBh$S7B7UcPH/iVTihLKH5sgIdBeS3O2X1StQsvv3cmbYd1  role network-admin
  ```

  In this case, the entered password should be **5$5$I4sapkBh$S7B7UcPH/iVTihLKH5sgIdBeS3O2X1StQsvv3cmbYd1**.

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

Cisco Nexus Dashboard Fabric Controller pushes feature lldp on the switches only for the easy fabric deployments, that is, for the eBGP-routed fabric or a VXLAN EVPN fabric.

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

  - External fabric in monitored or managed mode
  - LAN Classic fabric in monitored or managed Mode

- Backup/restore is only supported for Nexus devices on external fabrics.



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

- Beginning with Nexus Dashboard 4.3.1, the following Cisco Catalyst switches are supported for External fabric deployments:
  - Cisco Secure Router 8475-G2
  - Cisco Secure Router 8455-G2



The hardware module details for Cisco Secure Router 8475-G2 and Cisco Secure Router 8455-G2 are not displayed on the **Manage > Inventory > Switch > Hardware > Modules** page. To view this information, use the **show inventory** command from the device CLI.

  - Cisco Catalyst 9000 Series switches:
    - C9300L-24P-4G
    - C9300L-48P-4G
    - C9200CX-12P-2X2G-E



Traffic analytics compatibility mode is supported on C9300L-24P-4G and C9300L-48P-4G switches.

- Starting with Nexus Dashboard release 4.3.1, support is added for the Cisco N9164E-NS4-O switch running NX-OS release 10.6.2(n) or a later version.



The Cisco N9164E-NS4-O switch is limited to Layer-3 operations only. The following features are not supported on this switch:

- Layer-2: SVI, vPC (including pairing), and Port-channels.
- Traffic Management: Multicast, Port Access Control List (PACL), VLAN Access Control List (VACL), and VXLAN.
- Advanced Analysis: Connectivity Analysis, Netflow, and sFlow (resulting in no support for Traffic Analysis and Flow Telemetry).

Move an External Fabric Under an MSD Fabric

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

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

   The **Move Fabric** page displays. It contains a list of fabrics. The external fabric is displayed as a standalone fabric.

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

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

External Fabric Depiction in an MSD Fabric Topology

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



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

Add switches to the External fabric

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

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

   You can also add switches to an External fabric using these methods:

   - From **Manage > Fabrics**. Choose an External fabric and click **Inventory**. On the **Switches** tab, from the **Actions** drop-down list, click **Add switches** to add switches to the selected External fabric.
   - From **Topology**, right click on the External fabric and select **Add Switches**.

2. Determine how you want to add the switches to the External fabric.
   - Click **Discover** to discover new switches.
   - Click **Move Neighbor Switches** to add existing switches to the Fabric.

3. If you choose the **Discover** option:
   a. Enter the IP address (Seed IP) of the switch.
   b. In the **Authentication Protocol** field, from the drop-down list, select the appropriate protocol to add switches to the Fabric.
   c. Choose the device type from the **Device Type** drop-down list.
      - Choose **NX-OS** to discover a Cisco Nexus switch.
      - Choose **IOS XE** to discover a CSR device.
      - Choose **IOS XR** to discover an ASR device.
      - Choose **SONiC** to discover a switch that is running on the SONiC network operating system.

- Choose **Other** to discover non-Cisco devices.

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



- Config compliance is disabled for all non-Nexus devices except for Cisco CSR 1000v.
- Config compliance is not supported for switches running on the SONiC image. The **Replay Config** option is provided instead for the user to upload the JSON format configurations to merge or replace the running configurations on SONiC switches, as described below.

d. Enter the administrator username and password of the switch.

e. Click **Discover switches** at the bottom part of the screen.

The Scan Details section comes up shortly. Since the Max Hops field was populated with 2, the switch with the specified IP address and switches two hops from it are populated.

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

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

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

4. If you added switches that are running on the SONiC network operating system, upload the necessary JSON format configurations to merge or replace the running configurations on the SONiC switches.

   a. In the main **Switches** page, click **Actions > Maintenance > Replay config**.

   b. In the **Replay Config** page, upload the JSON format configurations file by clicking the box to navigate to the appropriate folder or by dragging the file to the box.

   c. Determine the **Replay mode** option that you want to use.

   - **Merge with existing running configs**
   - **Completely replace the existing running config**

   d. Click **Submit**.

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

   The selected switch is moved to the External fabric.

Switch Settings for External Fabrics

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

The options are:

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



Changing of switch role is allowed only before executing **Deploy Config**. **vPC Pairing** - Select a switch for vPC and then select its peer.

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

Manage Interfaces - Deploy configurations on the switch interfaces.

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

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

History - View per switch deployment history.

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

Deploy Config - Deploy per switch configurations.

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

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

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

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

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



If a switch in an external fabric does not accept default credentials, you should perform one of the following actions:
* Remove the switch in the external fabric from inventory, and then rediscover.
* LAN discovery uses both SNMP and SSH, so both passwords need to be the same. You need to change the SSH password to match the SNMP password on the switch. If SNMP authentication fails, discovery is stopped with authentication error. If SNMP authentication passes but SSH authentication fails, Nexus Dashboard Fabric Controller discovery continues, but the switch status shows a warning for the SSH error.

Discover new switches

To discover new switches, perform the following steps:

1. Power on the new switch in the external fabric after ensuring that it is cabled to the Nexus Dashboard Fabric Controller server.

   Boot the Cisco NX-OS and setup switch credentials.

2. Execute the **write**, **erase**, and **reload** commands on the switch.

   Choose **Yes** to both the CLI commands that prompt you to choose Yes or No.

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

   The **Edit Fabric** screen is displayed.

4. Click the **Bootstrap** tab and update the DHCP information.
5. Click **Save** at the bottom right part of the **Edit Fabric** screen to save the settings.
6. Double click on the Fabric to view the **Fabric Overview**.
7. On **Switches** tab, from the **Actions** drop-down list, select **Add Switches**.
8. Click the **POAP** tab.

   In an earlier step, the reload command was executed on the switch. When the switch restarts to reboot, Nexus Dashboard Fabric Controller retrieves the serial number, model number, and version from the switch and displays them on the Inventory Management along screen. Also, an option to add the management IP address, hostname, and password are made available. If the switch information is not retrieved, refresh the screen using the Refresh icon at the top right part of the screen.

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

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

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

   You can provision devices in advance.

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

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

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

10. (Optional) Use discovery credentials for discovering switches.
    a. Click the **Add Discovery Credentials** icon to enter the discovery credentials for switches.
    b. In the **Discovery Credentials** page, enter the discovery credentials such as discovery

username and password.

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

If the discovery credentials are not provided, Nexus Dashboard Fabric Controller uses the admin user and password to discover switches.



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

11. Click **Bootstrap** at the top right part of the screen.

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

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

12. Monitor and check the switch for POAP completion.
13. Click **Deploy Config** from the **Actions** drop-down list on the **Fabric Overview** screen to deploy pending configurations (such as template and interface configurations) onto the switches.

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

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

14. After the pending configurations are deployed, the **Progress** column displays 100% for all switches.
15. On the Topology screen, click **Refresh Topology** icon to view the update.

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

    The switch and the link are discovered in Nexus Dashboard Fabric Controller. Configurations are built based on various policies (such as fabric, topology, and switch generated policies). The

switch image (and other required) configurations are enabled on the switch.

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

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

17. On the Nexus Dashboard Fabric Controller UI, the discovered switches can be seen in the fabric topology.

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

    - vPC pairing.
    - Breakout interfaces

      Support for breakout interfaces is available for 9000 Series switches.

    - Port channels, and adding members to ports.



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

Adding Non-Nexus Devices to External Fabrics

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

- Cisco ASR 9000 Series Routers
- Cisco NCS 5500 Series Routers, IOS XR Release 6.5.3

From Cisco Nexus Dashboard Fabric Controller Release 12.1.1e, you can also add Cisco 8000 Series Routers to external fabrics both in managed mode and monitored mode.

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

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

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

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

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

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

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

Configuration Compliance in External Fabrics

With external fabrics, any Nexus switches, Cisco IOS-XE devices, Cisco IOS XR devices, and Arista can be imported into the fabric, and there is no restriction on the type of deployment. It can be LAN Classic, VXLAN, FabricPath, vPC, HSRP, etc. When switches are imported into an external fabric, the configuration on the switches is retained so that it is non-disruptive. Only basic policies such as the switch username and mgmt0 interface are created after a switch import.

In the external fabric, for any intent that is defined in the Nexus Dashboard Fabric Controller, configuration compliance (CC) ensures that this intent is present on the corresponding switch. If this intent is not present on the switch, CC reports an Out-of-Sync status. Additionally, there will be a Pending Config generated to push this intent to the switch to change the status to In-Sync. Any additional configuration that is on the switch but not in intent defined in Nexus Dashboard Fabric Controller, will be ignored by CC, as long as there is no conflict with anything in the intent.

When there is user-defined intent added on Nexus Dashboard Fabric Controller and the switch has additional configuration under the same top-level command, as mentioned earlier, CC will only ensure that the intent defined in Nexus Dashboard Fabric Controller is present on the switch. When this user defined intent on Nexus Dashboard Fabric Controller is deleted as a whole with the intention of removing it from the switch and the corresponding configuration exists on the switch, CC will report an Out-of-Sync status for the switch and will generate **Pending Config** to remove the config from the switch. This **Pending Config** includes the removal of the top-level command. This action leads to removal of the other out-of-band configurations made on the switch under this top-level command as well. If you choose to override this behavior, the recommendation is that, you create a freeform policy and add the relevant top-level command to the freeform policy.

Let us see this behavior with an example.

1. A **switch_freeform** policy defined by the user in Nexus Dashboard Fabric Controller and deployed to the switch.
2. Additional configuration exists under **router bgp** in **Running config** that does not exist in user-defined Nexus Dashboard Fabric Controller intent **Expected config**. Note that there is no **Pending Config** to remove the additional config that exists on the switch without a user defined intent on Nexus Dashboard Fabric Controller.
3. The **Pending Config** and the **Side-by-side Comparison** when the intent that was pushed earlier via Nexus Dashboard Fabric Controller is deleted from Nexus Dashboard Fabric Controller by deleting the switch_freeform policy that was created in the Step 1.
4. A **switch_freeform** policy with the top-level **router bgp** command needs to be created. This enables CC to generate the configuration needed to remove only the desired sub-config which was pushed from Nexus Dashboard Fabric Controller earlier.

5. The removed configuration is only the subset of the configuration that was pushed earlier from Nexus Dashboard Fabric Controller.

   For interfaces on the switch in the external fabric, Nexus Dashboard Fabric Controller either manages the entire interface or does not manage it at all. CC checks interfaces in the following ways:

   - For any interface, if there is a policy defined and associated with it, then this interface is considered as managed. All configurations associated with this interface must be defined in the associated interface policy. This is applicable for both logical and physical interfaces. Otherwise, CC removes any out-of-band updates made to the interface to change the status to **In-Sync**.
   - Interfaces created out-of-band (applies for logical interfaces such as port-channels, sub interfaces, SVIs, loopbacks, etc.), will be discovered by Nexus Dashboard Fabric Controller as part of the regular discovery process. However, since there is no intent for these interfaces, CC will not report an **Out-of-Sync** status for these interfaces.
   - For any interface, there can always be a monitor policy associated with it in Nexus Dashboard Fabric Controller. In this case, CC will ignore the interface's configuration when it reports the **In-Sync** or **Out-of-Sync** config compliance status.

Special Configuration CLIs Ignored for Configuration Compliance

The following configuration CLIs are ignored during configuration compliance checks:

- Any CLI having 'username' along with 'password'
- Any CLI that starts with 'snmp-server user'

Any CLIs that match the above will not show up in pending diffs and clicking Save & Deploy in the Fabric Builder window will not push such configurations to the switch. These CLIs will not show up in the Side-by-side Comparison window also.

To deploy such configuration CLIs, perform the following procedure:

1. Select **Manage > Fabrics**.

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

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

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

3. On the Policies tab, from the **Actions** drop-down list, select **Add policy**.
4. Add a Policy Template Instances (PTIs) with the required configuration CLIs using the **switch_freeform** template and click **Save**.
5. Select the created policy and select **Push Config** from the **Actions** drop-down list to deploy the configuration to the switch(es).

Managing Cisco IOS-XR Devices using NDFC

In general, workload requires communication with services outside of the data center domain in a data center fabric. This includes users accessing an application and services from the internet and

WAN. VXLAN EVPN fabrics with border devices are considered as a handoff for north-south connectivity. These border devices are in peer with IOS-XR routers, which are the backbone routers for WAN and internet connectivity.

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

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



For all non-Nexus devices, only the message-digest algorithm (MD5) protocol is supported for Simple Network Management Protocol version 3 (SNMPv3) authentication.



Beginning with NDFC 12.2.1, you do not need to configure the Simple Network Management Protocol (SNMP) for IOS-XR discovery of switches. NDFC uses Secure Shell (SSH) for IOS-XR device discovery.

Configuring IOS-XR as an Edge Router

To extend VRF Lite from a Cisco N9000 fabric with border devices for IOS-XR as the edge router, see the *VRF Lite Between Cisco N9000 Based Border and Non-Nexus Device* section for more information.

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

Configuring Non-Nexus Devices for Discovery

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

Configuring IOS-XE Devices for Discovery



In case of failure or issues configuring devices contact Cisco Technical Assistance Center (TAC). Before you discover the Cisco IOS-XE devices in Nexus Dashboard Fabric Controller, perform the following steps:

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

```
switch (config)# hostname <hostname>
switch (config)# ip domain name <domain_name>
switch (config)# crypto key generate rsa
switch (config)# ip ssh time-out 90
switch (config)# ip ssh version 2
```

```
switch (config)# line vty 1 4
switch (config-line)# transport input ssh
switch (config)# username admin privilege secret <password>
switch (config)# aaa new-model
switch (config)# aaa authentication login default local
switch (config)# aaa authorization exec default local none
```

2. Before you run SNMP command on the switch, ensure that the IP addresses, username and SNMP related configurations are defined on the switch. Run the following SNMP command on the switch console.

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

Configuring Arista Devices for Discovery

Enable Privilege Exec mode using the following command:

```
switch> enable
switch#

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

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

```
switch# configure terminal
switch (config)# username ndfc privilege 15 role network-admin secret passwd
snmp-server view _view_name_ SNMPv2 included
snmp-server view _view_name_ SNMPv3 included
```

```
snmp-server view _view_name_ default included
snmp-server view _view_name_ entity included
snmp-server view _view_name_ if included
snmp-server view _view_name_ iso included
snmp-server view _view_name_ lldp included
snmp-server view _view_name_ system included
snmp-server view sys-view default included
snmp-server view sys-view ifmib included
snmp-server view sys-view system included
snmp-server community private ro
snmp-server community public ro
snmp-server group _group_name_ v3 auth read _view_name_
snmp-server user username _group_name_ v3 auth md5 _password_ priv aes _password_
```



SNMP password should be same as the password for username.

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

Show Run Command

```
switch (config)# snmp-server engineID local f5717f444ca824448b00
snmp-server view _view_name_ SNMPv2 included
snmp-server view _view_name_ SNMPv3 included
snmp-server view _view_name_ default included
snmp-server view _view_name_ entity included
snmp-server view _view_name_ if included
snmp-server view _view_name_ iso included
snmp-server view _view_name_ lldp included
snmp-server view _view_name_ system included
snmp-server view sys-view default included
snmp-server view sys-view ifmib included
snmp-server view sys-view system included
snmp-server community private ro
snmp-server community public ro
snmp-server group _group_name_ v3 auth read _view_name_
snmp-server user _user_name__group_name_ v3 localized f5717f444ca824448b00 auth
md5 be2eca3fc858b62b2128a963a2b49373 priv aes
be2eca3fc858b62b2128a963a2b49373
!
spanning-tree mode mstp
!
service unsupported-transceiver labs f5047577
!
aaa authorization exec default local
```

```
!
no aaa root
!
username admin role network-admin secret sha512
$6$5ZKs/7.k2UxrWDg0$FOkdVQsBTnOquW/9AYx36YUBSPNLFdeuPIse9XgyHSdEOYXtPyT/
0sMUYYdkMffuljgn/d9rx/Do71XSbygSn/
username cvpadmin role network-admin secret sha512
$6$fLGFj/PUcuJT436i$Sj5G5c4y9cYjl/BZswjjmZW0J4npGrGqlyG3ZFk/ULza47Kz.d31q13jXA
7iHM677gwqQbFSH2/3oQEaHRq08.
username ndfc privilege 15 role network-admin secret sha512
$6$M48PNrCdg2ElTEdG$iiB880nvFQQlrWoZwOMzdt5EfkuClraNqtEMRS0TJUhNKCQnJN.VD
LFsLAmP7kQBo.C3ct4/.n.2eRlcP6hij/
```

Show SNMP View Command

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

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

Configuring and Verifying Cisco IOS-XR Devices for Discovery

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

```
switch# configure terminal
switch (config)# snmp-server viewview_namecisco included
snmp-server view _view_name_ mib-2 included
snmp-server group _group_name_ v3 auth read _view_name_ write _view_name_
```

```
snmp-server user _user_name__group_name_ v3 auth md5 password priv des56 password
SystemOwner
```

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

```
RP/0/RSP0/CPU0:ios(config)#snmp-server view view_name cisco included
RP/0/RSP0/CPU0:ios(config)#snmp-server view view_namemib-2 included
RP/0/RSP0/CPU0:ios(config)#snmp-server group group_name v3 auth read view_name
write view_name
RP/0/RSP0/CPU0:ios(config)#snmp-server user user_name_group_name_ v3 auth md5
password priv des56 passwordSystemOwner
RP/0/RSP0/CPU0:ios(config)#commit Day MMM DD HH:MM:SS Timezone
```

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

```
RP/0/RSP0/CPU0:ios(config)#
RP/0/RSP0/CPU0:ios(config)#show run snmp-server Day MMM DD HH:MM:SS Timezone
snmp-server user user_name group1 v3 auth md5 encrypted 10400B0F3A4640585851 priv
des56 encrypted 000A11103B0A59555B74 SystemOwner
snmp-server view _view_name_cisco included
snmp-server view _view_name_mib-2 included
snmp-server group group_name v3 auth read view_name write view_namev3 auth read
_view_name_ write _view_name_
```

Discovering Non-Nexus Devices in an External Fabric

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

Before you begin:

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

1. Navigate to **Manage > Fabrics**.
2. Double-click on an external fabric.

   The **Fabric Overview** page appears.
3. Click the **Switches** tab.
4. Choose **Add switches** in the **Actions** drop-down list.

   The **Add Switches - Fabric** page appears.

   The **Discover** radio button is selected in the **Switch Addition Mechanism** area.
5. Enter values for the following fields in the **Seed Switch Details** area.

Field	Description
Seed IP	Enter the IP address of the switch. You can discover more than one switch by providing the IP address range. For example: 10.10.10.40-60. The switches must be properly cabled and connected to the Nexus Dashboard Fabric Controller server and the switch status must be manageable.
Authentication Protocol	Choose the authentication protocol from the drop-down list. **MD5** is the default.
Device Type	• Choose **IOS XE** from the drop-down list for adding Cisco CSR 1000v, Cisco ASR 1000 series routers, Cisco Catalyst 8000 series switches, and Cisco Catalyst 9000 series switches. • Choose **IOS XR** from the drop-down list for adding ASR 9000 series routers, Cisco NCS 5500 and NCS 5001 series routers, IOS XR Release 6.5.3, or Cisco 8000 series routers.  To add Cisco IOS XR devices in managed mode, navigate to the **General Parameters** tab in the fabric settings and uncheck the **Fabric Monitor Mode** check box. • Choose **Other** from the drop-down list for adding non-Cisco devices, like Arista switches.
Username	Enter the username.
Password	Enter the password.
Set as individual device write credential	Specify if you want to set discovery/read credentials on individual devices.

6. Click **Discover Switches**.

The **Discovery Results** section appears with the switch details populated.



An error message appears if you try to discover a device that is already discovered.
Set the password of the device on the **LAN Credentials** page if the password is

not set. To navigate to the **LAN Credentials** page from the Cisco Nexus Dashboard Fabric Controller Web UI, choose **Admin > Switch Credentials > LAN Credentials Management**.

7. Check the check boxes next to the switches you want to add.
8. Click **Add Switches**.

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

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

   If you see the following error message after attempting to add the switch to the fabric:
   Error while creating the (Seed interface) intent for basic switch configurations. Please retry using config Save/Deploy.
   This might be because the permissions were not set properly for the switch before you tried to add it to the fabric. Set the permissions for the switch using the procedures in Configuring Non-Nexus Devices for Discovery, then try adding the switch to the fabric again.

9. Click **Close**.

   The fabric topology page appears with the switches.

10. (Optional) Click **Refresh** on the **Topology** page to view the latest topology view.
11. (Optional) Click **Fabric Overview**.

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

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

    After the discovery of the device:

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

    When a switch is in **Unreachable** discovery status, the last available information of the switch is retained in other columns. For example, if the switch was in **RUNNING** tracker status before it becomes unreachable, the value under the **Config Status** column for this switch will still be **RUNNING** despite the switch being in **Unreachable** discovery status.

What to do next:

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

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

Managing Non-Nexus Devices to External Fabrics

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



Configuration compliance is enabled for IOS-XE and IOS-XR switches, similar to the way the Nexus switches are handled in External Fabric. For more information, see Configuration Compliance in External Fabrics.
Nexus Dashboard Fabric Controller sends commit at the end of deployment for IOS-XR devices.
Nexus Dashboard Fabric Controller provides a few templates for IOS-XR devices. Use the **ios_xr_Ext_VRF_Lite_Jython.template** for IOS-XR switch to be an edge router to establish eBGP peering with border. This will create config for vrf, eBGP peering for the vrf and the sub-interface. Similarly, **ios_xe_Ext_VRF_Lite_Jython** can be used for IOS-XE switch to be an edge router to establish eBGP peering with border.

Understanding switches running SONiC

Support is now available for certain switches running SONiC (Software for Open Networking in Cloud) images in these Nexus Dashboard fabric types:

- AI Routed
- Routed
- External

Support is available for these switches running SONiC images:

- Cisco N9164E-NS4-O
- Cisco N93108TC-FX3 (N9K-C93108TC-FX3)

These sections provide more detailed information based on the fabric type.

- Switches running SONiC: On AI Routed and Routed fabrics
- Switches running SONiC: External fabrics
- Switches running SONiC: General information

Switches running SONiC: On AI Routed and Routed fabrics

For AI Routed and Routed fabrics, you will configure SONiC as the network operating system for the switches as part of the process of creating the fabric. See Creating Fabrics and Fabric Groups for more information.

When you choose the SONiC Network operating system in the **Advanced Configuration** mode when creating a fabric and you add the relevant switches to the fabric, the switches are bootstrapped through ZTP (Zero Touch Provisioning). During ZTP, the switch discovers onboarding parameters through DHCP and establishes an initial communication with Nexus Dashboard, which delivers the appropriate SONiC image and startup configuration.

New SONiC-specific fabric, interface, and network templates are also available as part of this feature.

Since the switches running the SONiC image are fully managed by Nexus Dashboard, when you add a configuration policy for the switch, only the leaf_bgp_asn policy is available in the **Add policy** page for you to choose. You would specify the leaf switch's ASN in the **Leaf BGP AS #** field in this page in the case where the **ASN auto allocation** option is disabled under **General Parameters** in fabric settings. See Working with Inventory in Your Nexus Dashboard LAN or IPFM Fabrics for more information.

The show commands, which are allowed to be run on SONiC switches, are listed in the sonic_allowed_show_clis template. The contents of that template are validated against when you choose **Actions > Maintenance > Show commands** and you choose sonic_show from the **Commands** pull-down list on the **Switch** show commands page.

Guidelines and limitations: Switches running SONiC on AI Routed and Routed fabrics

- You must open the Nexus Dashboard port 30022 for switches running SONiC images to function in Routed and AI Routed fabrics.
- Telemetry is always enabled on Routed and AI Routed fabrics with switches running SONiC images and cannot be disabled through the Nexus Dashboard GUI. It can be paused, which stops streaming but retains the configuration, and then resumed, which sends a full baseline then streams on-change updates. Telemetry can only be fully disabled by deleting the fabric.
- Per-packet dynamic load-balancing is supported on Cisco N9164E-NS4-O switches. Per-packet dynamic load-balancing is not supported on Cisco N93108TC-FX3 switches.
- AI job monitoring is not supported.

These are the additional guidelines and limitations for switches running SONiC.

- Deployment constraints: Switches running SONiC on AI Routed and Routed fabrics
- Unsupported features: Switches running SONiC on AI Routed and Routed fabrics

Deployment constraints: Switches running SONiC on AI Routed and Routed fabrics

- **Greenfield only**—Only greenfield deployments with IPv4 underlay and OOB management are supported.
- **Management access**—Devices can only be managed via OOB (eth0). Out-of-band configuration changes are not supported.
- **Multi-AS only**—Only Multi-AS mode is supported.
- **Port breakout**—Must be done through the Nexus Dashboard GUI workflow.
- **Default VRF**—Consistent with existing routed fabrics, only default VRF and external peering out of border switches in default VRF are supported.

Unsupported features: Switches running SONiC on AI Routed and Routed fabrics

The following features are **not** supported in this release:

Roles

In Routed and AI Routed fabrics with switches running SONiC images, only switch roles of spine, leaf and border are supported; all other switch roles are not supported.

Networking

- Port-channels, vPC, PVLAN
- HSRP/VRRP, BGP Authentication, BFD Authentication, MACsec



On the multi-attach of network screen, if a network is shown to be attached to short-formed interface eth0 for the SONiC switch, that interface actually refers to Ethernet0 rather than eth0 of the SONiC switch management interface. In other words:

- eth0 is the management port
- Ethernet0 to Ethernet47 are the front panel ports 1-48
- Ethernet48 is the front panel port 49. If this port is configured as the breakout port, then:
  - Ethernet48 to Ethernet51 are configured as the breakout port, and
  - Ethernet52 is the front panel port 50, and so on

Configuration

- Freeform configuration
- Banner and AAA configurations

Telemetry and Analytics

NetFlow, Flow Telemetry, Traffic Analytics, sFlow, Flow Rules (unavailable at the fabric level), Policy CAM Analysis, Delta Analysis, Traffic Analytics, Conformance, Connectivity Analysis, Energy Management, Log Collector, Bug Scan

Switches running SONiC: External fabrics

For External fabrics, you will configure SONiC as the network operating system for the switches as part of the process of adding switches to the External fabric and discovering those switches. See Editing External Fabric Settings for more information.

These are the additional guidelines and limitations for switches running SONiC.

- Deployment constraints: Switches running SONiC on External fabrics

Deployment constraints: Switches running SONiC on External fabrics

- **Config compliance and freeform**—For External fabrics, config compliance and freeform configs are not supported for switches running on the SONiC image. The **Replay config** option is provided instead for the user to upload the JSON format configurations to merge or replace the running configurations on the SONiC switches.

Switches running SONiC: General information

- General guidelines and limitations: Switches running SONiC
- General deployment constraints: Switches running SONiC
- General deployment model: Switches running SONiC

- Troubleshooting issues

General guidelines and limitations: Switches running SONiC

- A fabric that contains switches running SONiC images can only contain those switch types. You cannot have a mix of switches running SONiC images and switches with other types of images in the same fabric.
- For switches running on SONiC, you can only update the image on those switches using a SONiC image. In addition, you must choose **Disruptive** as the update type for SONiC images. See Managing Your Fabric Software for more information.
- Having the same subnet scope range between NX-OS fabrics and fabrics with switches running SONiC images is not supported in this release.

General deployment constraints: Switches running SONiC

- **Topology discovery**—LLDP-based only.

General deployment model: Switches running SONiC

Telemetry is only supported in a co-hosted deployment (the colocation model is not supported). This means fabrics with switches running SONiC images are not supported on virtual cluster (app 500G) cluster types or colocated setups.

Troubleshooting issues

To help facilitate troubleshooting, Nexus Dashboard provides **Begin** and **End** troubleshooting actions to create or update the credential of the user rescue-user with corresponding permissions on switches running SONiC images. You can then access the switches with the username rescue-user and execute commands for troubleshooting. Once you are finished with troubleshooting, you can choose the **End troubleshooting** action to remove rescue-user from the switches.

Creating a vPC Setup

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

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

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

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

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

   The **vPC Domain** and **vPC Peerlink** tabs appear. You must populate the fields in the tabs to

create the vPC setup. The description for each field is displayed at the extreme right.

Field	Description
vPC Domain tab	Enter the vPC domain details.
vPC+	If the switch is part of a FabricPath vPC + setup, enable this check box and enter the **FabricPath switch ID** field.
Configure VTEPs	Check this check box to enter the source loopback IP addresses for the two vPC peer VTEPs and the loopback interface secondary IP address for NVE configuration.
NVE interface	Enter the NVE interface. vPC pairing will configure only the source loopback interface. Use the freeform interface manager for additional configuration.
NVE loopback configuration	NVE loopback configuration: Enter the IP address with the mask. vPC pairing will only configure primary and secondary IP address for loopback interface. Use the freeform interface manager for additional configuration.
vPC Peerlink tab	Enter the vPC peer-link details.
Switch Port Mode	Choose **trunk** or **access** or **fabricpath**. If you select trunk, then corresponding fields (**Trunk Allowed VLANs** and **Native VLAN**) are enabled. If you select **access**, then the **Access VLAN** field is enabled. If you select **fabricpath**, then the trunk and access port related fields are disabled.

3. Click **Save**.

   The **vPC setup** is created.

   To update vPC setup details, do the following:

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

      The **vPC peer** dialog box comes up.

   b. Update the field(s) as needed.

      When you update a field, the **Unpair** icon changes to **Save**.

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

      After creating a vPC pair, you can view vPC details on the **vPC Overview** page.

Undeploying a vPC Setup

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

   The vPC peer screen comes up.

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

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

3. Click **Deploy Config**.
4. (Optional) Click the value under the **Recalculate Config** column.

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



Resync the fabric if it is out of sync.
When you unpair, only PTIs are deleted for following features, but the configuration is not cleared on the switch during **Deploy Config**: NVE configuration, LACP feature, fabricpath feature, nv overlay feature, loopback primary ID. In case of host vPCs, port channels and their member ports are not cleared. You can delete these port channels from the **Interfaces** window if required. You can continue using these features on the switch even after unpairing.
If you are migrating from fabricpath to VXLAN, you need to clear the configuration on the device before deploying the VXLAN configuration.

Precision Time Protocol for External Fabrics

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

The PTP feature is supported with Cisco N9000 Series cloud-scale switches, with NX-OS version 7.0(3)I7(1) or later. Warnings are displayed if there are non-cloud scale devices in the fabric, and PTP is not enabled. Examples of the cloud-scale devices are Cisco N93180YC-EX, Cisco N93180YC-FX, Cisco N93240YC-FX2, and Cisco N93360YC-FX2 switches. For more information, refer to https://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html.



Cisco N9348Y2C6D-SE1U and Cisco N9324C-SE1U switches do not support PTP. In brownfield deployments, onboarding fails if you add these switches to a fabric with PTP enabled.

Follow these steps to onboard switches:

1. Disable PTP at the fabric level before adding the switches.
2. Apply PTP configurations through free-form policies to the Cloud Scale switches that require PTP.
3. Enable Network Time Protocol (NTP) on the Cisco N9348Y2C6D-SE1U and Cisco N9324C-SE1U switches.



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

For more information, see the *Configuring PTP* chapter in *Cisco N9000 Series NX-OS System*

Management Configuration Guide and *Cisco Nexus Insights for Cisco User Guide*.

For External fabric deployments, you have to enable PTP globally, and also enable PTP on core-facing interfaces. The interfaces could be configured to the external PTP server like a VM or Linux-based machine. Therefore, the interface should be edited to have a connection with the grandmaster clock. For PTP and TTAG configurations to be operational on External fabrics, you must sync up of Switch Configs to Nexus Dashboard Fabric Controller using the **host_port_resync** policy. For more information, see the section "Sync up out-of-band switch interface configurations" in Working with Connectivity in Your Nexus Dashboard LAN Fabrics.

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

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

Here is the sample PTP configuration:

```
feature ptp

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

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

interface Ethernet1/59 -> Core facing interface
  ptp

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

The following guidelines are applicable for PTP:

- The PTP feature can be enabled in a fabric when all the switches in the fabric have Cisco NX-OS Release 7.0(3)I7(1) or a higher version. Otherwise, the following error message is displayed:

  ```
  PTP feature can be enabled in the fabric, when all the switches have
  NX-OS Release 7.0(3)I7(1) or higher version. Please upgrade switches to
  NX-OS Release 7.0(3)I7(1) or higher version to enable PTP in this fabric.
  ```

- For hardware telemetry support in NIR, the PTP configuration is a prerequisite.

- If you are adding a non-cloud scale device to an existing fabric which contains PTP configuration, the following warning is displayed:

  > TTAG is enabled fabric wide, when all devices are cloud-scale switches
  > so it cannot be enabled for newly added non cloud-scale device(s).

- If a fabric contains both cloud-scale and non-cloud scale devices, the following warning is displayed when you try to enable PTP:

  > TTAG is enabled fabric wide when all devices are cloud-scale switches
  > and is not enabled due to non cloud-scale device(s).

- TTAG configuration is generated for all the devices if host configuration sync up is performed on all the devices. TTAG configuration will not be generated for any newly added devices if host configuration sync up is not performed on all newly added devices.

  If the configuration is not synced, the following warning is displayed:

  > TTAG on interfaces with PTP feature can only be configured for cloud-scale devices.
  > It will not be enabled on any newly added switches due to the presence of non cloud-scale devices.

- PTP and TTAG configurations are deployed on host interfaces.
- PTP and TTAG Configurations are supported between switches in the same fabric (intra-fabric links). PTP is created for inter-fabric links, and TTAG is created for the inter-fabric link if the other fabric (Switch) is not managed by Nexus Dashboard Fabric Controller. Inter-fabric links do not support PTP or TTAG configurations if both fabrics are managed by Nexus Dashboard Fabric Controller.
- TTAG configuration is configured by default after the breakout. After the links are discovered and connected post breakout, perform **Deploy Config** to generate the correct configuration based on the type of port (host, intra-fabric link, or inter fabric link).

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