Editing Classic LAN Fabric Settings, Release 4.1.1

Table of Contents

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 4.1.1	Improved navigation and workflow when editing Classic LAN fabric settings.	Beginning with Nexus Dashboard 4.1.1, the navigation and workflow when editing Classic LAN fabric settings in Nexus Dashboard have been enhanced.

Editing Classic LAN fabric settings

A **Classic LAN** fabric is a type of fabric used for Access-Aggregation-Core Classic LAN architectures, based on Cisco best practices with Nexus 3000/7000/9000 switches. This fabric type supports the **Access** and **Aggregation** switch roles. Switches that serve as **Core** for these networks must be deployed in the **External and inter-fabric connectivity** fabric type.

When you first create a classic fabric using the procedures provided in Creating LAN and ACI 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 classic fabric.

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

   Manage > Fabrics

2. Locate the Classic LAN fabric that you want to edit.

   Classic LAN fabrics are shown with **Classic** in the **Type** column.

3. Click the circle next to the Classic LAN 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 Classic LAN fabric.

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

Field	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	Choose the location for the fabric.
VRF-Lite Protocol	Specifies the VRF-Lite Agg-Core/Edge or Collapsed Core-WAN peering protocol options. Options are: - EBGP - OSPF - None: Nexus Dashboard does not configure the peering protocol if the **None** option is selected. You must manually configure the peering protocol with this option, if necessary.
BGP ASN	Enter the BGP autonomous system number (ASN) for the fabric's spine switches.
License tier	Choose the licensing tier for the fabric: - **Essentials** - **Advantage** - **Premier** Click on the information icon (i) next to License tier to see what functionality is enabled for each license tier.
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 **Enable features** field above. Choose either **Out-of-band** or **In-band** for telemetry collection.
Telemetry streaming	This option becomes available if you choose to enable **Telemetry** in the **Enable features** field above. Choose either **IPv4** or **IPv6** for telemetry streaming.
Security domain	Choose the security domain for the fabric.

Fabric Management

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

- General Parameters
- Spanning Tree
- vPC
- Protocols
- Security
- Advanced
- Resources
- Manageability
- Bootstrap
- Configuration Backup
- 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
First Hop Redundancy Protocol	Specifies the FHRP protocol. Options are: - **none**: Select this option if you want Layer 2 only. - **hsrp** - **vrrp** - **vrrpv3**
Enable Performance Monitoring	Check the check box to enable performance monitoring. Ensure that you do not clear interface counters from the Command Line Interface of the switches. Clearing interface counters can cause the Performance Monitor to display incorrect data for traffic utilization. If you must clear the counters and the switch has both clear counters and clear counters snmp commands (not all switches have the clear counters snmp command), ensure that you run both the main and the SNMP commands simultaneously. For example, you must run the clear counters interface ethernet slot/port command followed by the clear counters interface ethernet slot/port snmp command. This can lead to a one time spike.

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

Spanning Tree

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

Field	Description
Spanning Tree Root Bridge Protocol	Specify the protocol to be used for configuring Root Bridge: Options are: • **rpvst+**: Rapid Per-VLAN Spanning Tree • **mst**: Multiple Spanning Tree • **unmanaged (default)**: STP Root not managed by Nexus Dashboard  Spanning Tree settings and bridge configurations are applicable at the Aggregation layer only.
Spanning Tree VLAN Range	Specify the VLAN range. For example: 1, 3-5, 7, 9-11 The default value is 1-3967. Applicable only for Aggregation devices.
MST Instance Range	Specify the MST instance range. For example: 0-3,5,7-9 The default value is 0. Applicable only for Aggregation devices.
Spanning Tree Bridge Priority	Specify the bridge priority for the spanning tree in increments of 4096. Applicable only for Aggregation devices.
Spanning Tree Hello Interval	Set the number of seconds between the generation of config spanning-tree Bridge Protocol Data Unit (BPDU). The default value is 2. Applicable only for Aggregation devices.
Spanning Tree Forward Delay	Set the number of seconds for the forward delay timer. The default value is 15. Applicable only for Aggregation devices.
Spanning Tree Max Age Interval	Set the maximum number of seconds the information in a spanning-tree Bridge Protocol Data Unit (BPDU) is valid. The default value is 20. Applicable only for Aggregation devices.
Spanning Tree Pathcost Method	Options are: • **short**: (default): Use 16-bit based values for default port path costs • **long**: Use 32-bit based values for default port path costs Applicable only for Aggregation devices.

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

vPC

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

Field	Description
vPC Auto Recovery Time	Specifies the vPC auto recovery time-out period in seconds. • Minimum value: 240 • Maximum value: 3600
vPC Delay Restore Time	Specifies the vPC delay restore period in seconds. • Minimum value: 1 • Maximum value: 3600
vPC Peer Link Port Channel ID	Specifies the Port Channel ID for a vPC Peer Link. Default value in this field is 500. • Minimum value: 1 • Maximum value: 4096
vPC IPv6 ND Synchronize	Enables IPv6 Neighbor Discovery synchronization between vPC switches. The check box is enabled by default. Uncheck the check box to disable the function.
vPC Domain Id Range	Specifies the vPC Domain Id range to use for new pairings.
vPC Layer-3 Peer-Router Option	Enables the Layer 3 device to form peering adjacency with both the peers. Configure this command in both the peers. If you configure this command only on one of the peers or you disable it on one peer, the operational state of layer 3 peer-router gets disabled. You get a notification when there is a change in the operational state.
Use Specific vPC/Port-Channel ID Range	Enable this check box to use a specific vPC/port-channel ID range for aggregation-access pairing.
vPC/Port-Channel ID Range	Specifies one vPC/port-channel ID range for auto-allocating vPC/port-channel IDs for aggregation-access pairing. The minimum allowed value is 1 and the maximum allowed value is 4096.

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

Protocols

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

Field	Description
OSPF Process Tag	This field becomes editable if you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. The OSPF Routing Process Tag. Maximum size is 20.
OSPF Area ID	This field becomes editable in these conditions: · If you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. · If you enter a value in the **OSPF Process Tag** field above. The **OSPF Area ID** in an IP address format.
OSPFv3 Process Tag	This field becomes editable if you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. The **OSPFv3 Routing Process Tag**. Maximum size is 20.
OSPFv3 Area ID	This field becomes editable in these conditions: · If you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. · If you enter a value in the **OSPFv3 Process Tag** field above. The **OSPFv3 Area ID** in an IP address format.
Enable BGP Authentication	This field becomes editable if you selected **ebgp** in the **Routing Protocol** field under the General Parameters tab. Select the check box to enable BGP authentication. Deselect the check box to disable it. If you enable this field, the **BGP Password Key Encryption Type** and **BGP Neighbor Password** fields are enabled.
BGP Password Key Encryption Type	This field becomes editable in these conditions: · If you selected **ebgp** in the **Routing Protocol** field under the General Parameters tab. · If you enabled the **Enable BGP Authentication** field above. Choose the 3 for 3DES encryption type, or 7 for Cisco encryption type.
BGP Neighbor Password	This field becomes editable in these conditions: · If you selected **ebgp** in the **Routing Protocol** field under the General Parameters tab. · If you enabled the **Enable BGP Authentication** field above. Enter the VRF Lite BGP neighbor password as a hex string.

Field	Description
Enable OSPF Authentication	This field becomes editable if you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. Select the check box to enable OSPF authentication. Deselect the check box to disable it. If you enable this field, the **OSPF Authentication Key ID** and **OSPF Authentication Key** fields get enabled.
OSPF Authentication Key ID	This field becomes editable in these conditions: · If you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. · If you enabled the **Enable OSPF Authentication** field above. The **Key ID** is populated.
OSPF Authentication Key	This field becomes editable in these conditions: · If you selected **ospf** in the **Routing Protocol** field under the General Parameters tab. · If you enabled the **Enable OSPF Authentication** field above. The OSPF authentication key must be the 3DES key from the switch. NOTE: Plain text passwords are not supported. Log in to the switch, retrieve the encrypted key and enter it in this field. For more information, see the *Retrieving the Authentication Key* section for details.

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

Security

The fields on the **Security** tab are described in the following table.

For more information on configuring data center interconnect (DCI) MACsec, see Connecting Two Fabrics with MACsec Using QKD.

Field	Description
Enable DCI MACsec	Check the check box to enable MACsec on DCI links.
Enable QKD	Check the check box to enable the QKD server for generating quantum keys for encryption. If you choose to not enable the **Enable QKD** option, Nexus Dashboard uses preshared keys provided by the user instead of using the QKD server to generate the keys. If you disable the **Enable QKD** option, all the fields pertaining to QKD are grayed out.

Field	Description
DCI MACsec Cipher Suite	Choose one of the following DCI MACsec cipher suites for the DCI MACsec policy: • **GCM-AES-128** • **GCM-AES-256** • **GCM-AES-XPN-128** • **GCM-AES-XPN-256** The default value is **GCM-AES-XPN-256**.
DCI MACsec Primary Key String	Specify a Cisco Type 7 encrypted octet string that is used for establishing the primary DCI MACsec session. For **AES_256_CMAC**, the key string length must be 130 and for **AES_128_CMAC**, the key string length must be 66. If these values are not specified correctly, an error is displayed when you save the fabric.  The default key lifetime is infinite.
DCI MACsec Primary Cryptographic Algorithm	Choose the cryptographic algorithm used for the primary key string. It can be **AES_128_CMAC** or **AES_256_CMAC**. The default value is **AES_128_CMAC**. You can configure a fallback key on the device to initiate a backup session if the primary session fails.
DCI MACsec Fallback Key String	Specify a Cisco Type 7 encrypted octet string that is used for establishing a fallback MACsec session. For **AES_256_CMAC**, the key string length must be 130 and for **AES_128_CMAC**, the key string length must be 66. If these values are not specified correctly, an error is displayed when you save the fabric.  This parameter is mandatory if the **Enable QKD** option is not selected.
DCI MACsec Fallback Cryptographic Algorithm	Choose the cryptographic algorithm used for the fallback key string. It can be **AES_128_CMAC** or **AES_256_CMAC**. The default value is **AES_128_CMAC**.
QKD Profile Name	Specify the crypto profile name. The maximum size is 63.
KME Server IP	Specify the IPv4 address for the Key Management Entity (KME) server.
KME Server Port Number	Specify the port number for the KME server.
Trustpoint Label	Specify the authentication type trustpoint label. The maximum size is 64.
Ignore Certificate	Enable this check box to skip verification of incoming certificates.

Field	Description
DCI MACsec Status Report Timer	Specify the DCI MACsec operational status periodic report timer in minutes.

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

Advanced

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

Field	Description
VRF Template	Specifies the VRF template for creating VRFs. These are pre-built, best practice templates for VRFs that are provided with Nexus Dashboard. You do not have to specify a template but one is automatically selected.
Network Template	Specifies the network template for creating networks. These are pre-built, best practice templates for networks that are provided with Nexus Dashboard. You do not have to specify a template but one is automatically selected.
Layer 2 Host Interface MTU	Specifies the MTU for the layer 2 host interface. This value should be an even number.
Unshut Host Interfaces by Default	Check this check box to unshut the host interfaces by default.
Power Supply Mode	Choose the appropriate power supply mode.
CoPP Profile	Choose the appropriate control plane policing (CoPP) profile 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.

Field	Description
Brownfield Network Name Format	Enter the format to be used to build the overlay network name during a brownfield import or migration. The network name should not contain any white spaces or special characters except underscore and hyphen. The network name must not be changed once the brownfield migration has been initiated. See the *Creating Networks for the Standalone Fabric* section for the naming convention of the network name. The syntax is [<string> | VLAN_ID] and the default value is **Auto_Net_VLANVLAN_ID**. When you create networks, the name is generated according to the syntax you specify. The following list describes the variables in the syntax: · **VLAN_ID**: Specifies the VLAN ID associated with the network. VLAN ID is specific to switches, hence Nexus Dashboard Fabric Controller picks the VLAN ID from one of the switches, where the network is found, randomly and use it in the name. We recommend not to use this unless the VLAN ID is consistent across the fabric. · **<string>**: This variable is optional and you can enter any number of alphanumeric characters that meet the network name guidelines. An example overlay network name: Site_VLAN1234 Ignore this field for greenfield deployments.
Enable CDP for Bootstrapped Switch	Enables CDP on management (mgmt0) interface for bootstrapped switch. By default, for bootstrapped switches, CDP is disabled on the mgmt0 interface.
Enable Tenant DHCP	Check the check box to enable feature dhcp and associated configurations globally on all switches in the fabric. This is a pre-requisite for support of DHCP for overlay networks that are part of the tenant VRFs. Ensure that **Enable Tenant DHCP** is enabled before enabling DHCP-related parameters in the overlay profiles.
Enable NX-API	Specifies enabling of NX-API on HTTPS.
NX-API HTTPS Port Number	Field becomes active if the **Enable NX-API** option is enabled. Enter the NX-API HTTPS port number. Default value is 443.

Field	Description
Enable HTTP NX-API	Specifies enabling of NX-API on HTTP. Enable this check box and the **Enable NX-API** check box to use HTTP. This check box is checked by default. If you uncheck this check box, the applications that use NX-API and supported by Cisco Nexus Dashboard Fabric Controller, such as Endpoint Locator (EPL), Layer 4-Layer 7 services (L4-L7 services), VXLAN OAM, and so on, start using 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	Field becomes active if the **Enable HTTP NX-API** option is enabled.Enter the NX-API HTTPS port number. Default value is 80.
Enable Strict Config Compliance	Enable the Strict Config Compliance feature by selecting this check box. It enables bi-directional compliance checks to flag additional configs in the running config that are not in the intent/expected config. By default, this feature is disabled.
Enable AAA IP Authorization	Enables AAA IP authorization, when IP Authorization is enabled in the remote authentication server. This is required to support Nexus Dashboard Fabric Controller in scenarios where customers have strict control of which IP addresses can have access to the switches.
Enable Nexus Dashboard as Trap Host	Select this check box to enable Nexus Dashboard Fabric Controller as an SNMP trap destination. Typically, for a native HA Nexus Dashboard Fabric Controller deployment, the eth1 VIP IP address will be configured as SNMP trap destination on the switches. By default, this check box is enabled.
Enable Agg/Access Auto Pairing	For back-to-back vPCs, enable this option to automatically pair aggregation and access devices based on topology.
Create Route-map fabric-rmap-redist-subnet	Enable this option to create a route map fabric-rmap-redist-subnet. This route-map matches tag 12345.
Greenfield Cleanup Option	Enable this field to clean the switch configuration without a reloads when PreserveConfig=no. Valid options are Enable or Disable.
Aggregation Freeform Config	Additional CLIs for all Aggregation devices as captured from show running configuration.
Access Freeform Config	Additional CLIs for all Access devices as captured from show running configuration.

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

Resources

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

Field	Description
Network VLAN Range	VLAN range for the per switch overlay network (min:2, max:4094).
Aggregation-Core/Aggregation-Edge Connectivity	Specify the option for the VRF Lite Aggregation-Core and Aggregation-Edge Router Inter-Fabric connection. Options are: - **Auto**: Automatically generates the VRF Lite configuration on the Aggregation and Core switches. This option is applicable only if you are using the Cisco Nexus 7000 or 9000 Series switches for the Core layer. - **Manual**: If you are using the Cisco Catalyst 9000 series switches or Cisco ASR 9000 Series Aggregation Services Routers for the Core layer, select **Manual** in this field. You must manually create a policy using the necessary policy provided to you through Nexus Dashboard. For more information, see VRF Lite.
VRF-Lite Subinterface dot1q Range	Specifies the per Aggregation dot1q Range for VRF Lite connectivity (min:2, max:4093).
Auto Generate VRF Lite Configuration on Aggregation and Core/Edge	Option that controls the automatic generation of the VRF Lite sub-interface and peering configurations on the Aggregation & Core/Edge devices. When this option is enabled, the automatically created VRF Lite links will have the 'Auto Generate Flag' enabled.
VRF Lite IP Version	Select the IP version for VRF Lite.Options are: - **IPv4_only** - **IPv6_only** - **IPv4_and_IPv6**
IPv4 VRF Subnet IP Range and **IPv4 VRF Subnet Mask Length**	The IPv4 address range to assign peer-to-peer Aggregation-Core connections, and peering between vPC Aggregation switches. Update the fields as needed. The values shown in your screen are automatically generated. If you want to update the IP address ranges or the VRF/Network VLAN ranges, ensure that it does not overlap with other ranges. You should only update one range of values at a time. If you want to update more than one range of values, do it in separate instances. For example, if you want to update Layer 2 and Layer 3 ranges, you should do the following: 1. Update the Layer 2 range and click **Save**. 2. Click the **Edit Fabric** option again, update the Layer 3 range and click **Save**.

Field	Description
IPv6 VRF Subnet IP Range and **IPv6 VRF Subnet Mask Length**	The IPv6 address range to assign peer-to-peer Aggregation-Core connections, and peering between vPC Aggregation switches. Update the fields as needed. The values shown in your screen are automatically generated. If you want to update the IP address ranges or the VRF/Network VLAN ranges, ensure that it does not overlap with other ranges. You should only update one range of values at a time. If you want to update more than one range of values, do it in separate instances. For example, if you want to update Layer 2 and Layer 3 ranges, you should do the following: 1. Update the Layer 2 range and click **Save**. 2. Click the **Edit Fabric** option again, update the Layer 3 range and click **Save**.
VRF Lite VLAN Range	VLAN range for Per VRF SVI Peering between Aggregation pairs (min:2, max:4094).

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

Manageability

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

Field	Description
DNS Server IPs	Specifies the comma separated list of IP addresses (v4/v6) of the DNS servers.
DNS Server VRFs	Specifies one VRF for all DNS servers or a comma separated list of VRFs, one per DNS server.
NTP Server IPs/Hostnames	Specifies a comma-separated list of IP addresses (IPv4/IPv6) or hostnames for the NTP server. Hostnames are limited to 80 characters in length and must not contain any whitespace or special characters, except for hyphens (-) and periods (.).
NTP Server VRFs	Specifies one VRF for all NTP servers or a comma separated list of VRFs, one per NTP server.
Syslog Server IPs/Hostnames	Specifies a comma-separated list of IP addresses (IPv4/IPv6) or hostnames for the Syslog server. Hostnames are limited to 199 characters in length and should not contain any whitespace or special characters, except for hyphens (-) and periods (.).
Syslog Server Severity	Specifies the comma separated list of syslog severity values, one per syslog server. The minimum value is 0 and the maximum value is 7. To specify a higher severity, enter a higher number.

Field	Description
Syslog Server VRFs	Specifies one VRF for all syslog servers or a comma separated list of VRFs, one per syslog server.
AAA Freeform Config	Specifies the AAA freeform configurations. If AAA configurations are specified in the fabric settings, **switch_freeform** PTI with source as **UNDERLAY_AAA** and description as **AAA Configurations** will be created.
Banner	Specifies the message of the day banner.

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

Bootstrap

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

Field	Description
Enable Bootstrap	Select this check box to enable the bootstrap feature. Bootstrap allows easy day-0 import and bring-up of new devices into an existing fabric. Bootstrap leverages the NX-OS POAP functionality. To add more switches and for POAP capability, chose check box for **Enable Bootstrap** and **Enable Local DHCP Server**. After you enable bootstrap, you can enable the DHCP server for automatic IP address assignment using one of the following methods: • External DHCP Server: Enter information about the external DHCP server in the **Switch Mgmt Default Gateway** and **Switch Mgmt IP Subnet Prefix** fields. • Local DHCP Server: Enable the **Local DHCP Server** check box and enter details for the remaining mandatory fields.
Enable Local DHCP Server	Select this check box to initiate enabling of automatic IP address assignment through the local DHCP server. When you select this check box, the **DHCP Scope Start Address** and **DHCP Scope End Address** fields become editable. If you do not select this check box, Nexus Dashboard Fabric Controller uses the remote or external DHCP server for automatic IP address assignment.

Field	Description
DHCP Version	Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the **Switch Mgmt IPv6 Subnet Prefix** field is disabled. If you select DHCPv6, the **Switch Mgmt IP Subnet Prefix** is disabled.  Cisco Nexus 9000 and 3000 Series Switches support IPv6 POAP only when switches are either Layer-2 adjacent (eth1 or out-of-band subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other than /64 are not supported.
DHCP Scope Start Address and **DHCP Scope End Address**	Specifies the first and last IP addresses of the IP address range to be used for the switch out of band POAP.
Switch Mgmt Default Gateway	Specifies the default gateway for the management VRF on the switch.
Switch Mgmt IP Subnet Prefix	Specifies the prefix for the Mgmt0 interface on the switch. The prefix should be between 8 and 30. *DHCP scope and management default gateway IP address specification*: If you specify the management default gateway IP address 10.0.1.1 and subnet mask 24, ensure that the DHCP scope is within the specified subnet, between 10.0.1.2 and 10.0.1.254.
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** For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24
Enable AAA Config	Select this check box to include AAA configurations from the Manageability tab as part of the device start-up config post bootstrap.
Bootstrap Freeform Config	(Optional) Enter additional commands as needed. For example, if you require some additional configurations to be pushed to the device and be available post device bootstrap, they can be captured in this field, to save the desired intent. After the devices boot up, they will contain the configuration defined in the **Bootstrap Freeform Config** field. Copy-paste the running-config to a **freeform config** field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. For more information, see Enabling Freeform Configurations on Fabric Switches.

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

Configuration Backup

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

Field	Description
Hourly Fabric Backup	Select the check box to enable an hourly backup of fabric configurations and the intent. The hourly backups are triggered during the first 10 minutes of the hour.
Scheduled Fabric Backup	Check the check box to enable a daily backup. This backup tracks changes in running configurations on the fabric devices that are not tracked by configuration compliance.
Scheduled Time	Specify the scheduled backup time in a 24-hour format. This field is enabled if you check the **Scheduled Fabric Backup** check box. Select both the check boxes to enable both back up processes. The backup process is initiated after you click **Save**. The scheduled backups are triggered exactly at the time you specify with a delay of up to two minutes. The scheduled backups are triggered regardless of the configuration deployment status. You can enter the number of fabric backups that will be retained on Nexus Dashboard in **Admin > System Settings > Fabric management > Advanced settings > LAN-Fabric > Maximum Backups per Fabric**. Default is 2. The number of archived files that can be retained is set in the **# Number of archived files per device to be retained:** field in the **Server Properties** page. Note: To trigger an immediate backup, do the following: 1. Choose **Overview > Topology**. 2. Click within the specific fabric box. The fabric topology screen comes up. 3. Right-click on a switch within the fabric, then select **Preview Config**. 4. On the **Preview Config** page for this fabric, click **Re-Sync All**. You can also initiate the fabric backup on the fabric topology page. Click **Backup Now** in the **Actions** pane.

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

Flow Monitor

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

automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Enable Netflow	Check this check box to enable Netflow on Aggregation devices for this fabric. By default, Netflow is disabled. On Enable, NetFlow configuration will be applied to all Aggregation devices 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 Nexus Dashboard, see the "Configuring Netflow support" section in Creating LAN and ACI Fabrics and Fabric Groups.

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

Field	Description
Exporter Name	Specifies the name of the exporter.
IP	Specifies the IP address of the exporter.
VRF	Specifies the VRF over which the exporter is routed.
Source Interface	Enter the source interface name.
UDP Port	Specifies the UDP port over which the netflow data is exported.

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

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

Field	Description
Record Name	Specifies the name of the record.
Record Template	Specifies the template for the record. Enter one of the record templates names.

The following two record templates are available for use. You can create custom netflow record templates. Custom record templates saved in the template library are available for use here.

- **netflow_ipv4_record** - to use the IPv4 record template.
- **netflow_l2_record** - to use the Layer 2 record template.
    - **Is Layer2 Record** - Check this check box if the record is for Layer2 netflow.

Click **Save** to configure the report. Click **Cancel** to discard. You can also choose an existing record

and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

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

Field	Description
Monitor Name	Specifies the name of the monitor.
Record Name	Specifies the name of the record for the monitor.
Exporter1 Name	Specifies the name of the exporter for the netflow monitor.
Exporter2 Name (optional)	Specifies the name of the secondary exporter for the netflow monitor.

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

In the **Netflow Sampler** area, click **Actions > Add** to add one or more Netflow samplers. These are optional fields and are applicable only when there are N7K aggregation switches in the fabric. The fields on this screen are:

Field	Description
Sampler Name	Specifies the name of the sampler.
Number of Samples	Specifies the number of samples.
Number of Packets in Each Sampling	Specifies the number of packets in each sampling.

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

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

Telemetry

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.

   

   NOTE: 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**.

Add the onboarded NAS to Nexus Dashboard

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

1. Navigate to the Fabrics page:

    Manage > Fabrics

2. Choose the fabric with the telemetry feature enabled.
3. Choose **Actions > Edit Fabric Settings**.
4. Click **Telemetry**.
5. Click the **NAS** tab in the **Telemetry** window.
6. Make the necessary configurations in the **General settings** area.
   a. Enter the name in the **Name** field.
   b. In the **NAS server** field, choose the NAS server added to Nexus Dashboard from the drop-down list.
7. In the **Collection settings** area, choose the flow from the **Flows** drop-down list.
   - 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.

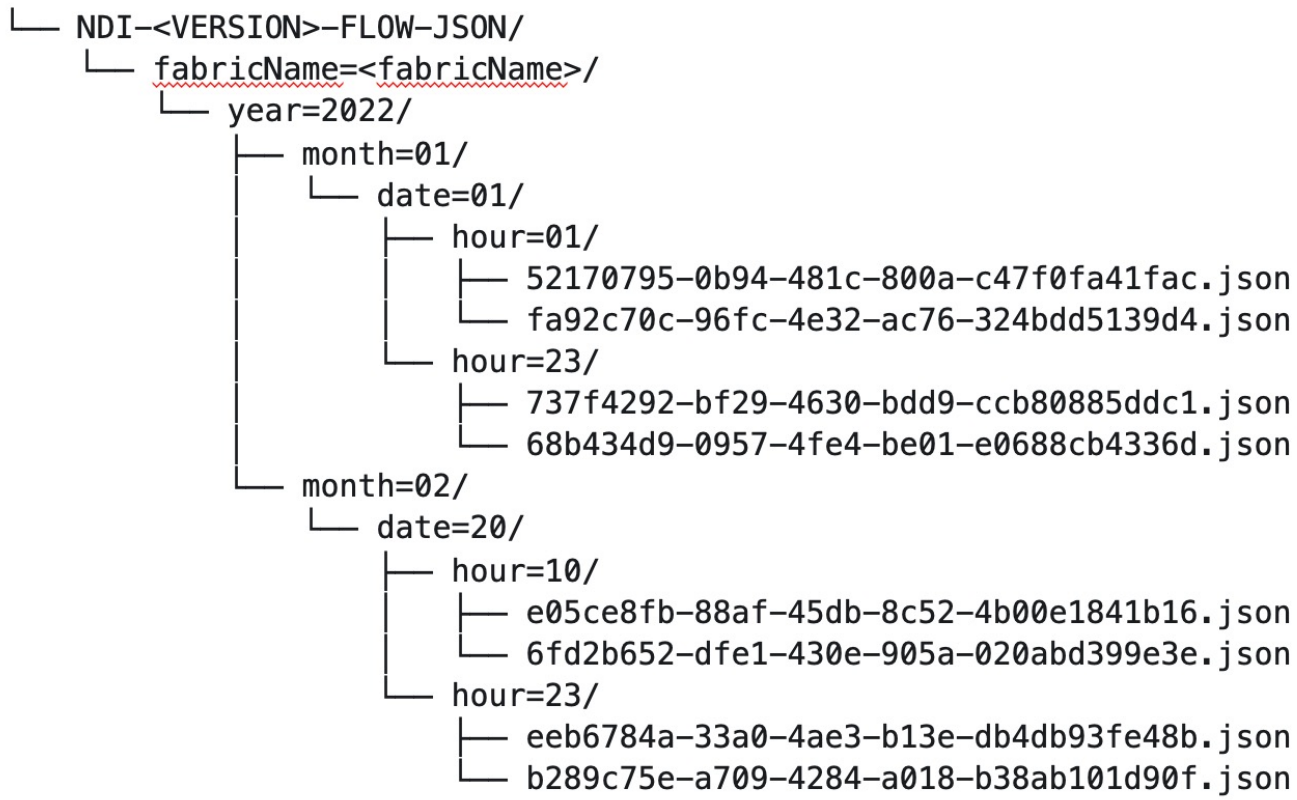

Navigate to **Analyze** > **Flows** to view the flows that will be exported.

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" :" U

DP" }

{" 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" ," egress Vrf" :" vrf1" ," ingressTenant" :" FSV2" ," egressTenant" :" FSV2" ," protocol" :" ANY-HOST" ," srcEpg" :" ext-epg" ," dstEpg" :" ext-epg1" ," latencyMax" :0," ingressVif" :" eth1/15" ," ingressVni" :0," latency" :0," ingressNodes" :" Leaf1-2" ," ingressVlan" :0," ingressByteCount" :104681600," ingressPktCount" :817825," 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.
- Nexus Dashboard supports Kafka export for Flow anomalies. However, Kafka export is not currently supported for Flow Event anomalies.

Flow telemetry guidelines and limitations for NX-OS fabrics

- Ensure that you have configured NTP and enabled PTP in Nexus Dashboard. See Cisco Nexus Dashboard Deployment Guide and Precision Time Protocol (PTP) for Cisco Nexus Dashboard Insights for more information. You are responsible for configuring the switches with external NTP servers.
- In the **Edit Flow** page, you can enable all three telemetry types. sFlow is most restrictive, Netflow has some more capability, and Flow Telemetry has the most capability. We recommend that you enable Flow Telemetry if it is available for your configuration. If Flow Telemetry is not available, then use Netflow. If Netflow is not available, use sFlow.
- If there are multiple Nexus Dashboard clusters onboarded to Nexus Dashboard, partial paths will be generated for each fabric.
- If you manually configure the fabric to use with Nexus Dashboard and Flow Telemetry support, the Flows Exporter port changes from 30000 to 5640. To prevent a breakage of the Flows Exporter, adjust the automation.
- Nexus Dashboard supports Kafka export for Flow anomalies. However, Kafka export is not currently supported for Flow Event anomalies.
- Flow telemetry is supported in -FX3 platform switches for the following NX-OS versions:
  - 9.3(7) and later
  - 10.1(2) and later
  - Flow telemetry is not supported in -FX3 platform switches for NX-OS version 10.1(1).
- Interface based Flow Telemetry is only supported on modular chassis with -FX land -GX line cards on physical ports and port-channels rules.
- If interface-based Flow Telemetry is pushed from Nexus Dashboard for **Classic LAN** and **External Connectivity Network** fabrics, perform the following steps:
  - Choose the fabric.

- Choose **Policies** > **Action** > **Add policy** > **Select all** > **Choose template** > host_port_resync and click **Save**.
- In the Fabric Overview page, choose **Actions** > **Recalculate and deploy**.

- For VXLAN fabrics, interface-based Flow Telemetry is not supported on switch links between spine switch and leaf switch.
- If you want to use the default VRF instance for flow telemetry, you must create the VRF instance with a name of " default" in lowercase. Do not enter the name with any capital letters.
- Flow telemetry is not supported in classic LAN topologies with 2-level VPC access layers.
- If you want to enable Flow Telemetry, ensure that there are no pre-existing Netflow configurations on the switches. If there are any pre-existing configurations, the switch configuration may fail.

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

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

      Manage > Fabrics

   b. In the table showing all of the Nexus Dashboard fabrics that you have already created, locate the LAN or IPFM fabric where you want to configure telemetry settings.
   c. Single-click on that fabric.

      The **Overview** page for that fabric appears.

   d. Click **Actions > Edit Fabric Settings**.

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

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

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

   f. Click the **Telemetry** tab to access the telemetry settings for this fabric.
2. Click the **Flow collection** tab in the **Telemetry** window.
3. In the **Mode** area, click **Flow telemetry**.
4. In the **Flow collections rules** area, determine what sort of flow collection rule that you want to add.
   - VRF
   - Physical interface
   - Port channel

VRF

To add a VRF rule:

1. Click the **VRF** tab.

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

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

2. Add a new rule by clicking **Create flow collection rule**.
   a. In the **General** area, complete the following:

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

ii. The VRF field is disabled. The flow rule applies to all the VRF instances.

iii. In the **Flow Properties** area, select the protocol for which you intend to monitor the flow traffic.

iv. Enter the source and destination IP addresses. Enter the source and destination port.

v. Click **Save**.

Physical interface

To add a physical interface rule:

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

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

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

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

Port channel

To add a port channel rule:

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

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

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

2. Add a new rule by clicking **Create flow collection rule**.
   a. In the **General** area, enter the name of the rule in the **Rule Name** field.
      i. Select the **Enabled** check box to enable the status. If you enable the status, the rule will take effect. Otherwise, the rule will be removed from the switches.

ii. In the **Flow Properties** area, select the protocol for which you intend to monitor the flow traffic.

iii. Enter the source and destination IP addresses. Enter the source and destination port.

iv. From the drop-down list, select an interface. You can add more than one row (node+interface combination) by clicking **Add Interfaces**. However, within the rule, a node can appear only once. Configuration is rejected if more than one node is added.

v. Click **Save**.

3. Click **Done**.

Monitor the subnet for flow telemetry

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

Follow these steps to monitor the subnet for flow telemetry.

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

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

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

Understanding Netflow

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

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

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

The components of a flow monitoring setup are as follows:

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

Supported interfaces for Netflow

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

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

Understanding Netflow types

You can use these Netflow types.

Full Netflow

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

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

Sampled Netflow

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

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

Netflow guidelines and limitations

- In Cisco Nexus 9000 series switches, Netflow supports a small subset of the published export fields in the RFC.
- Netflow is captured only on the ingress port of a flow as only the ingress switch exports the flow. Netflow cannot be captured on fabric ports.
- You must configure persistent IP addresses under the cluster configuration, including 7 IP addresses in the same subnet as the data network.

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** > **Flow collection**. In the **Flow Collection per Fabric** table, click **View** in the **Collector List** column.

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

Netflow guidelines and limitations for NX-OS fabrics

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

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

Configure Netflow

Follow these steps to configure Netflow.

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

Understanding sFlow

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

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

The following interfaces are supported for sFlow:

Supported interfaces for sFlow

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

Guidelines and limitations for sFlow

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

Configure sFlow telemetry

Prerequisites

Follow these steps to configure sFlow telemetry.

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

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.

To configure email at the system settings level, see [Add email configuration].

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.

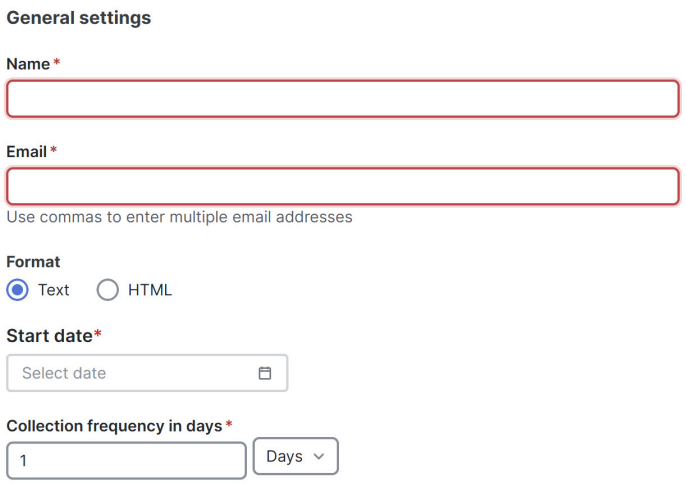

2. Follow these steps to configure **Collection Settings**.
   a. In the **Mode** field, choose one of the following modes.
      - **Basic**—displays the severity levels for anomalies and advisories.
      - **Advanced**—displays the categories and severity levels for anomalies and advisories.
   b. 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.

Field	Description
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 and Identifying Advisories in Your Nexus Dashboard.

Advisories

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

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

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

  For more information on advisory levels, see Detecting Anomalies and Identifying Advisories in Your Nexus Dashboard.

Statistics

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

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

- Interfaces
- Protocol
- Resource Allocation
- Environmental
- Endpoints

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.

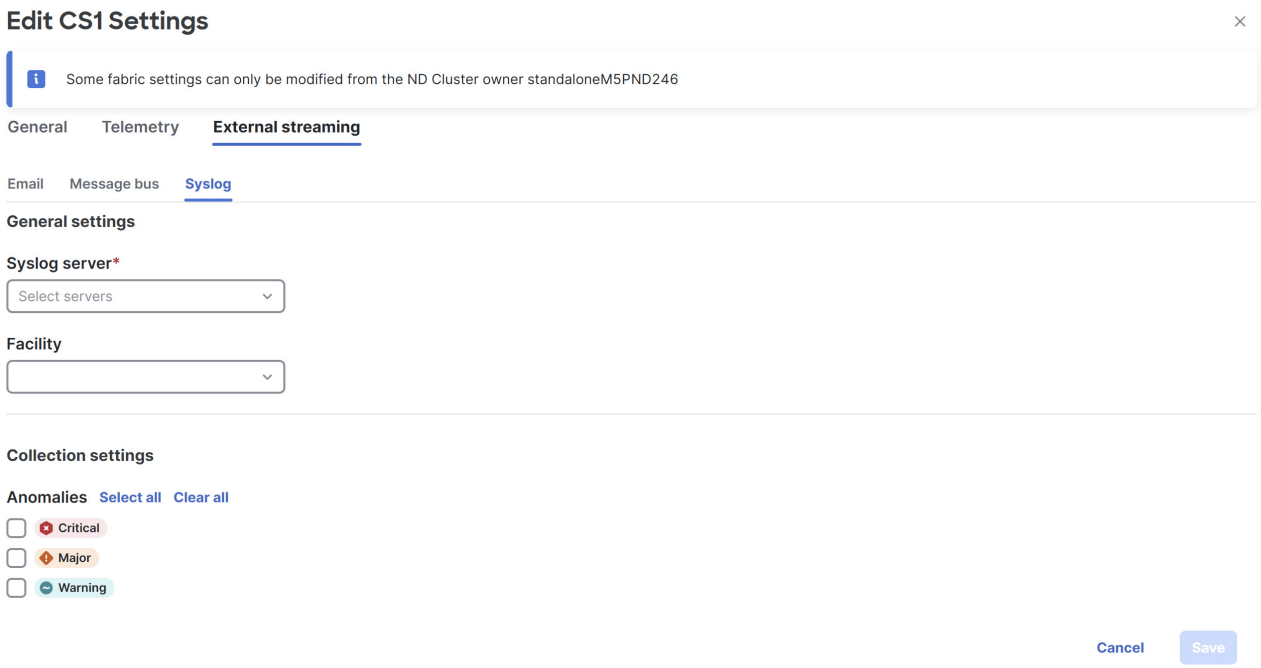

6. Make the necessary configurations in the **General settings** area.
   a. In the **Syslog server** drop down list, choose a syslog server.

      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 a classic fabric.

About aggregation-access pairing in a Classic LAN fabric

Nexus Dashboard has a one-click vPC feature for automatically detecting and pairing aggregation and access switches for optimal traffic engineering. By default, the auto aggregation-access pairing option is enabled, which means that after you perform a **Recalculate and deploy** operation, Nexus Dashboard automatically detects the connectivity between the aggregation and the access switches and generates the appropriate configurations based on the detected supported topologies. The configurations include vPC domains that Nexus Dashboard automatically pushes to the paired aggregation and access switches. The links between these aggregation-access pairs are bundled into a common vPC logical construct.

Workflow for configuring aggregation-access pairing

1. Create a Classic LAN fabric. For more information, see Creating LAN and ACI Fabrics and Fabric Groups.
2. Discover the switches in the fabric. For more information, see the section "Adding Switches to a Fabric" in Configuring Switches for LAN and IPFM Fabrics.
3. Add the switches using a bootstrap. For more information, see the section "Adding Switches Using Bootstrap Mechanism" in Configuring Switches for LAN and IPFM Fabrics.
4. Define the roles for the aggregation and access switches. For more information, see the section "Assigning Switch Roles" in Configuring Switches for LAN and IPFM Fabrics.
5. Configure the vPC pairing. For more information, see the section "Creating a vPC Setup" in Configuring Switches for LAN and IPFM Fabrics.
6. Configure aggregation-access pairings. For more information, see Create aggregation-access pairings.
7. Recalculate and deploy.

Create aggregation-access pairings

1. Perform the following procedure to configure an aggregation and an access switch, where aggregation switches are connected to access switches through a port channel.
2. Add an aggregation and an access switch to a Classic LAN fabric and set the role as either **Access** or **Aggregation** depending on the type of switch.

   With a Classic LAN fabric, Nexus Dashboard supports a minimum of two aggregation switches and the aggregation switches must be in a vPC pair.
3. On the fabric **Overview** page for the Classic LAN fabric, click **Inventory**, then choose an aggregation switch.
4. Click **Actions > Access Pairing**.

The **Access Pairing** page displays the aggregation switches on the top and a list of potential pairing access switches below the aggregation switches. The pairing status is displayed in the **Details** column.

5. Choose **Action > Edit Pairing**.

   The **aggregation** page displays.

6. Check the **Enable <switch-name> Pairing as Access Pairing** check box to pair the switches.
7. Click **Save**.
8. On the **Fabric Overview** page, click **Actions > Recalculate and deploy**.
9. After the configuration deployment is completed on the **Deploy Configuration** page, click **Close**.

Unpair aggregation-access switches

1. Uncheck the **Enable <switch-name> Pairing as Access Pairing** check box to unpair the switches.

   You cannot unpair an aggregation-access pair if overlays are attached.

2. On the fabric **Overview** page, click **Actions > Recalculate and deploy** to complete the unpairing operation.

Specifying a vPC/Port-Channel ID Range and Providing Custom vPC/Port-Channel IDs for Aggregation-Access Pairing

With this feature, you can:

- Configure a specific vPC/port-channel ID range for aggregation-access pairing by enabling the **Use Specific vPC/Port-Channel ID Range** field. Nexus Dashboard then displays the **vPC/Port-Channel ID Range** field with the recommended vPC/port-channel ID range.
- Edit vPC/port-channel IDs for paired switches.

Configure Fabric Settings for Specifying a vPC/Port-Channel ID Range for Aggregation-Access Pairing

1. On the **Fabric Overview** page, create an **Enhanced Classic LAN** fabric. For more information, see [Creating an Enhanced Classic LAN Fabric].
2. Click on the **vPC** tab.
3. Check the **Use Specific vPC/Port-Channel ID Range** check box to use a specific vPC/port-channel ID range for aggregation-access pairing.

   The **vPC/Port-Channel ID Range** field displays the recommended values.

   The recommended values are from 1-499.

   

   You can increase the existing range or add more ranges if the values are

exhausted.

4. Specify a range for the **vPC/Port-Channel ID Range** field if you do not want to use the recommended values.
5. Click **Save**.

   The new range applies to the new pairing.

Edit the Aggregation or the Access vPC/Port Channel IDs

1. On the **Fabric Overview > Switches** page, choose the aggregation switch you want to edit and click **Actions > Access Pairing**.

   The **Access Pairing** page appears with a horizontal bar of the paired aggregation switches.

2. Click **Edit Pairing** under the **Action** column.

   The access-aggregation paired switches page displays.

   The **Enable <switch-name> Pairing as Access Pairing** check box is checked due to auto aggregation-access pairing.

3. Click the arrow on the right-hand column of the page to view the fields.
4. Modify the access or the aggregation vPC/port-channel IDs if you want to change the values.
5. Click **Save**.

If you have overlays attached to the paired switches, you cannot change the vPC/port-channel IDs.

6. Navigate to the **Fabric Overview > Switches** page and click **Actions > Recalculate and deploy**.

   The **Deploy Configuration** page displays with the list of aggregation switches.

   After successful deployment, the **Fabric Status** column displays as **In-Sync**.

First Published: 2025-01-31
Last Modified: 2025-01-31