Working with System Settings, Release 4.3.1

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 4.3.1	Kerberos proxy authentication	Beginning with Nexus Dashboard 4.3.1, you can configure Kerberos-based authentication for outbound HTTP and HTTPS proxy traffic. As a Kerberos client, Nexus Dashboard retrieves tickets from the configured Key Distribution Center (KDC) to authenticate with the proxy. For more information, see Proxy configuration.
Nexus Dashboard 4.3.1	Support for CyberArk credential store for SAN fabrics	Beginning with Nexus Dashboard 4.3.1, support for CyberArk credential store is extended to SAN fabrics. For more information, see AAA remote authentication passthrough for SAN fabrics.

Navigate to System Settings

To navigate to the **System Settings** page, click **Admin > System Settings**.

These tabs are available under the **System Settings** page:

- General
- Fabric management
- Multi-cluster connectivity
- Flow collection

General

The **General** page under **System Settings** has these areas:

- Cluster details
- DNS
- NTP
- Proxy configuration
- Routes
- External pools
- Advanced settings
- Remote storage
- Remote streaming servers
- History and log settings
- Stream system anomalies
- Metadata
- Message bus configuration
- Email

Cluster details

Use the **Cluster details** area to view these cluster details:

- Name
- App subnet
- Service subnet
- App subnet IPv6
- Service subnet IPv6

DNS

Use the **DNS** area to configure DNS for the system.

Guidelines and limitations: DNS

- The DNS server must work for both UDP and TCP in order to successfully add it in the **Providers** area, as described in Configure DNS.

Configure DNS

To configure DNS:

1. Navigate to System Settings.

2. Click **General**.
3. In the **DNS** area, click **Edit**.
4. Add DNS search domains.
    a. Under **Search domains**, click **+ Add DNS search domain**.
    b. Enter the DNS search domain in the field, then click the checkmark on that row.
    c. Repeat these steps to add additional DNS search domains.
5. Add DNS providers.
    a. Under **Providers**, click **+ Add DNS provider**.
    b. Enter the DNS provider in the field, then click the checkmark on that row.
    c. Repeat these steps to add additional DNS providers.
6. Click **Save**.

NTP

Use the **NTP** area to configure NTP for the system.

1. Navigate to System Settings.
2. Click **General**.
3. In the **NTP** area, click **Edit**.
4. Add NTP ID keys.
    a. Under **ID**, click **+ Add key**.
    b. Enter the necessary information or enable options in these fields:
        - ID
        - Key
        - Trusted checkbox
        - Authentication type
    c. Click the checkmark on that row when you have completed the NTP ID configuration.
    d. Repeat these steps to add additional NTP ID keys.
5. Add NTP hostname/IP addresses.
    a. Under **NTP hostname/IP address**, click **+ Add NTP hostname/IP address**.
    b. Enter the necessary information or enable options in these fields:
        - NTP hostname/IP address
        - Key
        - Preferred checkbox
    c. Click the checkmark on that row when you have completed the NTP hostname/IP address configuration.
    d. Repeat these steps to add additional NTP hostname/IP address configurations.
6. Click **Save**.

Proxy configuration

Use the **Proxy configuration** area to configure a proxy server.

Add a proxy server

In certain deployment scenarios, you may have to access the Internet through a proxy.

Note that Nexus Dashboard uses two main route tables—one for the Management network and one for the Data network—and by default, it will use the routing table of the originating IP address. In other words, Nexus Dashboard will attempt to reach the proxy from the routing table of the POD/Service that is trying to use the proxy.

For example, if you configure a proxy and establish Intersight connectivity from your Nexus Dashboard and then attempt to configure the AppDynamics integration from the Insights service running in the cluster, you may get an error stating that the AppDynamics host is not reachable. This happens because the proxy is only accessible from the management interface, so in such cases you also need to add a management network route for the proxy IP address.

Beginning with Nexus Dashboard release 4.3.1, you can configure Kerberos-based authentication for outbound proxy traffic. In earlier releases, Nexus Dashboard supported only proxy configurations that used no authentication or username and password authentication.

With Kerberos proxy authentication, Nexus Dashboard acts as a Kerberos client. Nexus Dashboard authenticates to your organization's Key Distribution Center (KDC), obtains the required Kerberos tickets, and uses those tickets to authenticate to the configured proxy for all outbound HTTP and HTTPS traffic.

Kerberos proxy authentication applies only to outbound Nexus Dashboard traffic to the configured HTTP or HTTPS proxy. Nexus Dashboard does not use Kerberos to authenticate users or to authenticate directly to destination websites or services.

Before you begin

- Ensure that the proxy is configured to support Kerberos authentication.
- Ensure that Nexus Dashboard can reach the Kerberos Key Distribution Center (KDC).
- Ensure that time is synchronized between Nexus Dashboard and the KDC.

Guidelines and limitations: Kerberos proxy authentication

- Nexus Dashboard supports only one Kerberos realm for proxy authentication. You cannot configure multiple realms or domain-realm mappings.
- Nexus Dashboard manages runtime Kerberos tickets. These tickets are not user-configurable.
- Kerberos relies on time-sensitive tickets. Time skew between Nexus Dashboard and the KDC can cause authentication failures. Ensure that both systems use the same NTP source.
- The configured KDC and proxy server must be compatible with the Nexus Dashboard cluster networking mode:
  - In an IPv4-only deployment, Nexus Dashboard rejects IPv6-only values.
  - In an IPv6-only deployment, Nexus Dashboard rejects IPv4-only values.

    - In a dual-stack deployment, Nexus Dashboard allows either IPv4 or IPv6 values.
- The Kerberos password is stored securely in the Credential Management Service (CMS). Runtime Kerberos tickets (TGT and TGS) are transient and are not backed up, restored, or preserved across a restart or upgrade. After a restart, upgrade, or restore, Nexus Dashboard automatically reacquires fresh tickets by using the stored credentials.
- If Kerberos ticket acquisition fails after you save the configuration, Nexus Dashboard raises a system anomaly. For information about viewing system anomalies, see View platform and system alerts in global anomalies table.

Follow these steps to add a proxy server.

1. Navigate to **Admin > System Settings > General**.
2. In the **Proxy configuration** area, click **Edit**.
3. Click **+ Add proxy server** in the **Proxy Configuration** page.

   To configure Kerberos authentication for an existing proxy server, click the edit icon for the proxy entry. To replace an existing proxy configuration, delete an existing proxy entry and then click **+ Add proxy server**.
4. From the **Type** dropdown, select the type of traffic that you want to be proxied.

   If you choose **HTTP** and **HTTPS**, Nexus Dashboard applies the same proxy configuration to both HTTP and HTTPS traffic.
5. In the **Server** field, provide the full address for the proxy server, including the port if required.

   For example http://proxy.company.com:80.

   You can specify the proxy server as a hostname, IPv4 address, or IPv6 address. If you specify the proxy server as an IP address, the Kerberos environment must support an IP-based service principal name (SPN) for the proxy.
6. Choose the required authentication type:
    - None
    - Credentials
    - Kerberos
7. If you choose Credentials, provide the **Username** and **Password**.
8. If you choose **Kerberos**, add the Kerberos realm.
    a. In the **Realm** area, click **+ Add realm**.
    b. In the **Add Realm** page, provide the realm details:

Field	Description
Name	Kerberos realm name.
KDC	Hostname, IPv4 address, or IPv6 address of the Kerberos KDC.
Username	Username that Nexus Dashboard uses to authenticate to the KDC.
Password	Password for the Kerberos username.

c. Click **Save**.

9. (Optional) In the **Client principal name** field, enter the Kerberos client principal.

   If you do not provide a value, Nexus Dashboard generates the client principal from the username and realm.

10. (Optional) In the **Service principal name** field, enter the Kerberos service principal for the proxy.

    If you do not provide a value, Nexus Dashboard generates the service principal from the proxy server hostname and realm.

11. Click **Save**.

    Nexus Dashboard saves the Kerberos proxy configuration only after it successfully validates the KDC and acquires the required Kerberos tickets. After validation succeeds, Nexus Dashboard displays the validation status and the time of the last successful validation.

12. (Optional) Click **+ Add Ignore Host** to provide any hosts that bypass the proxy.

    You can add one or more hosts that the cluster communicates with directly, bypassing the proxy.

13. Click **Save**.

    Nexus Dashboard records proxy configuration changes in the audit logs.

Routes

Configure management and data network routes through **Routes**. You might configure management or data network routes for several reasons, such as:

- When using inband POAP, you will have to configure the appropriate data network routes for reachability to the switch inband interfaces. For more information, see Configuring Inband Management and Out-of-Band PnP.
- Configuring your Git IP route in order to import a template from Git. For more information, see Managing the Template Library for LAN Fabrics.

To configure management and data network routes:

1. Navigate to System Settings.
2. Click **General**.
3. In the **Routes** area, click **Edit**.

   The **Routes** page appears, with the configured management and data network routes displayed.

   - To configure additional management network routes, in the **Management network routes** area, click **+ Add management network routes**. Enter the IP address for the management network route, then click the checkmark in that row.
   - To configure additional data network routes, in the **Data network routes** area, click **+ Add data network routes**. Enter the IP address for the data network route, then click the checkmark in that row.

External pools

Configure persistent IP addresses through **External pools**.

1. Navigate to System Settings.
2. Click **General**.
3. In the **External Pools** area, click **Edit**.

   The **External Service Pools** page appears.

4. In the **Persistent data IPs** area, click **+ Add IP Address** and enter the necessary persistent IP addresses. Click the checkmark for each IP address that you enter.

Change persistent IP address

You can change the persistent IP addresses that are assigned for mandatory pods, such as POAP-SCP and SNMP traps.

To change the persistent IP address, perform the following steps:

1. On the Nexus Dashboard Web UI, navigate to **Admin > System Settings > Fabric Management**.
2. Under **Advanced Settings**, click **Admin**.
3. In the **LAN Device Management Connectivity** field, change **Management** to **Data** or vice versa.

   Changing the option results in a migration of SNMP and POAP-SCP pods to the persistent IP addresses associated with **External Service Pool** on Nexus Dashboard associated with the new **LAN Device Management Connectivity** option. After the completion of this process, the following message is displayed:

   Some features have been updated. Reload the page to see latest changes.

   Click **Reload the page**.

4. On the Nexus Dashboard Web UI, navigate to **Admin > System Settings > General**.
5. In the **External pools** card, click **Edit** to change the required IP addresses for **Persistent management IPs** or **Persistent data IPs**.
6. Navigate back to **Admin > System Settings > Fabric Management > Advanced Settings > Admin**, then change the option in **LAN Device Management Connectivity** drop-down list to its initial selection.

   Restoring this option to initial settings results in migration of the SNMP and POAP-SCP pods to use the updated persistent IP address from the appropriate external Service IP pool.

Advanced settings

You can configure these areas under **Advanced settings**:

- Fabric snapshot creation
- Display advanced settings and options for TAC support

Fabric snapshot creation

To enable the snapshot feature at the system level.

1. Navigate to **Admin > System Settings**.
2. With the **General** tab selected, locate the **Advanced Settings** area.
3. Determine if the **Fabric snapshot creation** feature is enabled or not.
    - If you see **Enabled** under the **Fabric snapshot creation** field, this feature has already been enabled.
    - If you see **Disabled** under the **Fabric snapshot creation** field:
        a. Click **Edit** in the Advanced Settings area.

        The **Advanced settings** slide-in pane appears.

        b. Click the checkbox next to **Enable fabric snapshot creation** to enable this feature, then click **Save**.

        You will now see **Enabled** under the **Fabric snapshot creation** field.

See the section "Onboard snapshot LAN fabrics" in Creating Fabrics and Fabric Groups for more information.

Display advanced settings and options for TAC support

To display advanced settings and options for TAC support.

1. Navigate to **Admin > System Settings**.
2. With the **General** tab selected, locate the **Advanced Settings** area.
3. Determine if the **Display advanced settings and options for TAC support** feature is enabled or not.
    - If you see **Enabled** under the **Display advanced settings and options for TAC support** field, this feature has already been enabled.
    - If you see **Disabled** under the **Display advanced settings and options for TAC support** field:
        a. Click **Edit** in the Advanced Settings area.

        The **Advanced settings** slide-in pane appears.

        b. Click the checkbox next to **Display advanced settings and options for TAC support** to enable this feature, then click **Save**.

        You will now see **Enabled** under the **Display advanced settings and options for TAC support** field.

Remote storage

The remote storage location information is referenced by any feature that uses remote storage location, including the unified backup and restore.

1. In the Nexus Dashboard GUI, navigate to **Admin > System Settings**.

   The **General** tab is chosen by default.

2. Locate the **Remote storage** area in **General**.
   - If you do not have any remote storage locations already created, you will see the message No remote storage added displayed on the page.
   - If you have remote storage locations already created, you'll see those remote storage locations listed with the following values:

Field	Description
Name	The name of the remote storage location.
Description	A description of the remote storage location, if necessary.
IP Address	The IP address of the remote storage location.
Protocol	The remote storage location type: • NAS Storage • SFTP
Status	The status of the remote storage location.

3. If there are no remote storage locations created yet, click **Edit** in the **Remote storage** area.

   The **Remote storage** page appears.

4. Click **+ Add Remote Storage Locations** to create a remote storage location.

   The **Create Remote Storage Location** page appears.

5. In the **Create Remote Storage Location** page, enter the necessary information to configure the remote storage location.

Field	Description
Name	Enter the name of the remote storage location.
Description	(Optional) Enter a description of the remote storage location.
Remote Storage Location Type	Choose **SFTP/SCP Server** as the remote storage location type.  As mentioned earlier, any feature that uses a remote storage location references the remote storage location information, not just unified backup and restore. Even though the **Create Remote Storage Location** shows **NAS Storage** as an option, it is not supported with the unified backup and restore feature.

   Enter the necessary information when choosing the **SFTP/SCP Server** option in the **Remote Storage Location Type** field.

Field	Description
Protocol	Choose the protocol to use for the remote storage location file transfer: • SFTP • SCP
Hostname or IP Address	Enter the hostname or IP address of the remote storage location.
Default Path	Enter the path to the directory where the backup file is to be saved on the remote server. • The path can be an absolute path, which would start with a slash character (/), such as: /backups/multisite • Or the path can be a path relative to your home directory, such as: Users/backups/multisite
Remote Port	Enter the remote port for the remote host location. This field is pre-populated with a default value of 22.
Authorization Type	Choose the authorization type: • Password • SSH Public Types • CyberArk
Username	Enter the authorization username.
Password	Available if you chose **Password** in the **Authorization Type** field above. Enter the authorization password.
SSH Key	The **SSH Key** and **Passphrase** fields are available if you chose **SSH Public Types** in the **Authorization Type** field. To use SSH keys, follow these steps: 1. Generate the private/public key pairs, with or without a passphrase. 2. Authorize the generated public key on the remote storage location. 3. Enter the private key in the **SSH Key** field. 4. If you used a passphrase in step 1, enter the passphrase in the **Passphrase** field.
Passphrase	
Credential Store key	The **Credential Store key** field is available if you choose **CyberArk** in the **Authorization Type** field. You will see the **CyberArk** tab only if you configured system certificate and mapped to CyberArk feature. For CA certificate information, refer to Managing Certificates in your Nexus Dashboard. For credential store details, refer to Configuring Users, Roles, and Security for LAN and AAA remote authentication passthrough for SAN fabrics for SAN.

6. Click **Save**.

   You are returned to the **Remote storage** page with the newly-created remote storage location listed in the table.

   - To edit a remote storage location entry, click on the ellipsis (**...**) at the end of the row in the table for that remote storage location and click **Edit**.
   - To delete a remote storage location entry, click on the ellipsis (**...**) at the end of the row in the table for that remote storage location and click **Delete**.

7. Click Save in the **Remote storage** page.

   You are returned to the **System Settings/General** page.

Add network attached storage to export flow records

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

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

Add NAS to Nexus Dashboard

Follow these steps to add NAS to Nexus Dashboard.

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

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

   d. In the **Hostname** field, enter the IP address of the Network Attached Storage.
   e. In the **Port** field, enter the port number of the Network Attached Storage.
   f. In the **Export path** field, enter the export path.

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

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

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

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

i. Click **Save**.

> In a three-node Nexus Dashboard cluster deployment, NAS backup operations require that all cluster nodes are in an operational and healthy state. If any node is down, unreachable, or not in a fully functional state, NAS backup operations do not proceed.
>
> This ensures cluster consistency and data integrity during backup operations. Verify cluster health and confirm that all three nodes are up and synchronized before initiating or scheduling NAS backups.

Add the onboarded NAS to Nexus Dashboard

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

1. Navigate to the Fabrics page:

   Manage > Fabrics

2. Choose the fabric with the telemetry feature enabled.
3. Choose **Actions > Edit Fabric Settings**.
4. Click **Telemetry**.
5. Click the **NAS** tab in the **Telemetry** page.
6. Make the necessary configurations in the **General settings** area.
   a. Enter the name in the **Name** field.
   b. In the **NAS server** field, choose the NAS server added to Nexus Dashboard from the drop-down list.
7. In the **Collection settings** area, choose the flow from the **Flows** drop-down list.
   - In Base mode, only 5-tuple data for the flow record is exported.
   - In Full mode, the entire data for the flow record is exported.
8. Click **Save**.

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

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

[nas directory] | *nas-directory.jpg*

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

Each flow record is written as a line delimited JSON.

JSON output file format for a flow record in base mode

{" fabricName" :" myapic" ," terminalTs" :1688537547433," originTs" :1688537530376," srcIp"

:" 2000:201:1:1::1" ," dstIp" :" 2000:201:1:1::3" ," srcPort" :1231," dstPort" :1232," ingressVrf" :" vrf1" ," egressVrf" :" vrf1" ," ingressTenant" :" FSV1" ," egressTenant" :" FSV1" ," protocol" :" UDP" }

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

JSON output file format for a flow record in full mode

{" fabricName" :" myapic" ," terminalTs" :1688538023562," originTs" :1688538010527," srcIp" :" 201.1.1.121" ," dstIp" :" 201.1.1.127" ," srcPort" :0," dstPort" :0," ingressVrf" :" vrf1" ," egressVrf" :" vrf1" ," ingressTenant" :" FSV2" ," egressTenant" :" FSV2" ," protocol" :" ANY-HOST" ," srcEpg" :" ext-epg" ," dstEpg" :" ext-epg1" ," latencyMax" :0," ingressVif" :" eth1/15" ," ingressVni" :0," latency" :0," ingressNodes" :" Leaf1-2" ," ingressVlan" :0," ingressByteCount" :104681600," ingressPktCount" :817825," ingressBurst" :0," ingressBurstMax" :34768," egressNodes" :" Leaf1-2" ," egressVif" :" po4" ," egressVni" :0," egressVlan" :0," egressByteCount" :104681600," egressPktCount" :817825," egressBurst" :0," egressBurstMax" :34768," dropPktCount" :0," dropByteCount" :0," dropCode" :" " ," dropScore" :0," moveScore" :0," latencyScore" :0," burstScore" :0," anomalyScore" :0," hashCollision" :false," dropNodes" :" []" ," nodeNames" :" [\" Leaf1-2\"]" ," nodeIngressVifs" :" [\" Leaf1-2,eth1/15\"]" ," nodeEgressVifs" :" [\" Leaf1-2,po4\"]“ ," srcMoveCount" :0," dstMoveCount" :0," moveCount" :0," prexmit" :0," rtoOutside" :false," events" :" [[\\\" 1688538010527,Leaf1-2,0,3,1,no,no,eth1/15,,po4,po4,,,,,0,64,0,,,,,,,,\\\"]]" }

Remote streaming servers

Use the Remote streaming servers area to add remote streaming servers.

Add remote streaming servers

You might need remote streaming servers to be configured when you deal with various areas in Nexus Dashboard, such as tech support, anomalies, and history and logs.

Follow these steps to add remote streaming servers.

1. Choose **Admin > System Settings**.

   The **General** tab is chosen by default.

2. Locate the **Remote streaming servers** tile and click **Edit**.

   A table showing any already-configured remote streaming servers appears.

3. Click **Add Server**.

The **Add Server** dialog box opens.

4. Determine the service where you want to add remote streaming servers.
   - Splunk
   - Syslog
   - SNMP
   - Webhook

Splunk

1. In the **Protocol** field, choose **HTTP** or **HTTPS** based on the Splunk URL that you use.

   For example:

   - Choose **HTTP**, if your login URL is http://splunk.mycompany.com:8000.
   - Choose HTTPS, if your login URL is https://mycompany.splunkcloud.com/, https://mycompany.splunkcloud.com:443, or https://splunk.mycompany.com.

   > Before enabling HTTP or HTTPS, you must upload the CA certificate for the Splunk destination host to Nexus Dashboard.

2. Enter a name for the Splunk remote streaming server in the **Name** field.
3. In the **Hostname/IP Address**, enter the hostname or IP address.
4. In the **Port** field, enter the HTTP event collector (HEC) port number.

   The HEC port number for Splunk deployments over HTTP or HTTPS protocols can be one of the following:

   - 8088 on Splunk Cloud free trials
   - 443 by default on Splunk Cloud Platform instances

5. In the **Token** field, enter the HEC token.

   For more information on Splunk deployment guidelines, see the Set up and use HTTP Event Collector in Splunk Web.

6. In the **Index** field, enter the necessary index information for the Splunk remote streaming server.

   > The index details on the Nexus Dashboard should match with the Splunk index.

7. Click **Save** in the **Add Server** page.

   You are returned to the **Remote streaming server** page.

8. In the **Remote streaming server** page, click **Add Server** to add another remote streaming server or click **Save** to save the configured remote streaming servers.

Syslog

1. In the **Protocol** field, choose either **TCP** or **UDP**.

2. Enter a name for the Syslog remote streaming server in the **Name** field.
3. In the **Hostname/IP Address** and **Port** fields, enter the hostname or IP address and port details.
4. Check the box next to the **TLS** field to enable this option.

   Before enabling TLS, you must upload the CA certificate for the Syslog destination host to Nexus Dashboard. See Managing Certificates in your Nexus Dashboard for more information.

5. Click **Save** in the **Add Server** page.

   You are returned to the **Remote streaming server** page.

6. In the **Remote streaming server** page, click **Add Server** to add another remote streaming server or click **Save** to save the configured remote streaming servers.

The CA certificate you upload is used to sign the certificate that Splunk and Syslog uses. The certificate for Splunk and Syslog should contain the IP address or DNS name in the Subject Alternative Name (SAN) section.

SNMP

Previously, Nexus Dashboard anomaly notifications were limited to Kafka or syslog exports. Beginning with Nexus Dashboard 4.3.1, the SNMP anomaly export feature enables network operations to seamlessly integrate Nexus Dashboard anomaly data into the existing SNMP trap collectors and Network Management Systems (NMS). This allows for centralized anomaly collection, simplifies monitoring, and helps meet compliance policies. You can selectively export system anomalies based on severity (Critical, Major, Minor, and Warning), ensuring focus on the most relevant alerts. This provides a more versatile and compliant method for consuming Nexus Dashboard anomaly data, especially for environments heavily reliant on SNMP.

1. Enter a name for the SNMP remote streaming server in the **Name** field.
2. In the **Hostname/IP Address** and **Port** fields, enter the SNMP server's hostname or IP address and port details.
3. Choose the SNMP **Version**:
   - **v2c**: Choose **v2c**.
     - In the **Community string** field, enter the community string name.
   - **v3**: Choose **v3**.
     - Enter the SNMPv3 engine ID if it is different from the SNMP server engine ID.
     - Choose your option from the **Authentication / Privacy** drop-down list. If you choose **Authentication / No privacy** from the drop-down list, then choose from **Authentication algorithm** drop-down list for the authentication protocol (MD5, SHA, SHA2-224, SHA2-256, SHA2-384, SHA2-512), enter the **Password** and **Username**.
     - If you choose **Authentication / privacy** from the drop-down list, then choose from **Authentication algorithm** drop-down list for the authentication protocol (MD5, SHA, SHA2-224, SHA2-256, SHA2-384, SHA2-512), enter the **Password**. Choose from **Encryption algorithm** drop-down list, enter the **Password** and **Username**.
     - If you choose **No authentication / No privacy** from the drop-down list, enter the **Username**.

4. Click **Save** in the **Add Server** page.

   You are returned to the **Remote streaming server** page.

5. In the **Remote streaming server** page, click **Add Server** to add another remote streaming server or click **Save** to save the configured remote streaming servers.

- Multiple SNMP servers of different versions (SNMPv2c and SNMPv3) can be configured and coexist simultaneously.
- To stream anomalies to an external destination, configure Nexus Dashboard for system anomaly streaming. For more information, refer to Stream system anomalies.

Webhook

Starting with Nexus Dashboard 4.2.1, Nexus Dashboard supports streaming event data to external webhook servers, in addition to existing options such as **Syslog**, **SNMP**, and **Splunk**. Webhook export feature enables network operations to seamlessly integrate Nexus Dashboard anomaly data to remote Webhook endpoints. This allows for centralized anomaly collection, simplifies monitoring, and helps meet compliance policies.

You can selectively export fabric-level anomalies and advisories based on severity (Critical, Major, Minor, Warning), ensuring focus on the most relevant alerts. This provides a more versatile and compliant method for consuming Nexus Dashboard anomaly data.

This feature adds **Remote Streaming capability** via webhook endpoints, allowing administrators to receive real-time event notifications.

- In Nexus Dashboard 4.2.1, this feature supports exporting only Fabric anomalies and advisories. System anomalies are not supported for Webhook streaming in this release.
- Webhook streaming is supported only for IPv4. IPv6 support is not available in this release.

Enable Webhook as a remote streaming server

You can enable Webhook as a remote streaming server and configure Webhook endpoints to receive real-time event notifications.

Follow these steps to enable Webhook remote streaming.

1. Follow the steps described in the Add remote streaming servers section, and click on **Webhook**.
2. Enter a name for the **Webhook remote streaming server** in the **Name** field.
3. (Optional) Enter a description for the Webhook remote streaming server in the **Description** field.
4. Enter the URL in the URL field using the format **http://{hostname/IP}:{port}/{endpoint}.**

   (Optional) Choose **Proxy**. If enabled, the URL will follow the Nexus Dashboard proxy.

5. (Optional) Enter the header name in the **Header name** field.
6. (Optional) Enter the header value in the **Header value** field.

Click **Add** to add more headers.

7. Click **Save** in the **Add Server** page.

Add Webhook streaming for a fabric

After adding the **Webhook** configuration in **Admin > System Settings > General > Remote streaming servers**, you can enable Webhook streaming for fabric-level anomalies and advisories.

Follow these steps to add Webhook streaming for a fabric.

1. Navigate to **Manage > Fabrics**.
2. Click the radio button next to the fabric that you want to edit, then click **Actions > Edit fabric settings**.

   The **Edit *fabric_name* settings** page appears.

3. Click the **External streaming** tab and then choose **Webhook**.
4. In the **Destination** drop-down list, choose the Webhooks you want to configure for your fabric.
5. Choose the **Severity of Anomalies** and **Severity of Advisories**.

   You can check the check box for each severity type (Critical, Major, Minor, Warning) or click **Select all** to choose all severity types.

6. Click **Save**.

History and log settings

Use the **History and log settings** area to enable history and log setting and set the maximum audit record retention duration for your Nexus Dashboard.

To enable history and log setting and set the maximum audit record retention duration for your Nexus Dashboard:

1. Go to **Admin > System Settings > General**.
2. Under **History and logs settings**, click **Edit**.

   The **History and logs settings** dialog box opens.

3. Check the box in the **Enable audit logs streaming** field to enable that feature, then choose a remote streaming server from the drop-down list.

   See Add remote streaming servers for more information on adding remote streaming servers.

4. In the **Maximum Number Record Retention duration** field, enter the number of months as the duration that you want to have the records retained.

   Valid entries are from 1-12. The default duration is 12 months.

5. In the **Maximum Number of Audit Records** field, enter the maximum number of audit records that you want to retain.

Valid entries are from 0-2147483647. The default value for maximum number of audit records is 100K for virtual cluster setup and 500K for a physical cluster. NOTE: When a network reachability issue occurs between Nexus Dashboard and Splunk or Syslog, logs generated during the outage are typically not replayed after the connectivity is restored. Instead, log streaming resumes from the point of recovery, resulting in the loss of any log data produced during the downtime. This means that any log data generated during the downtime is lost and not backfilled or retransmitted after the network is re-established.

Stream system anomalies

Use the **Stream system anomalies** area to configure destination syslog servers for streaming anomaly records. For more information, see Detecting Anomalies in Your Nexus Dashboard.

1. Navigate to System Settings.
2. Click **General**.
3. In the **Stream system anomalies** area, click **Edit**.
4. Select a remote streaming server (Syslog, SNMP, or Warning) from the **Remote streaming servers** drop-down list.
5. In the **Anomalies** area, choose the anomaly level, or click Select all to choose all of the anomaly levels.

   Anomaly levels include critical, major, and minor.
6. Click **Save**.

Metadata

To view or configure metadata information:

1. Navigate to **Admin > System Settings**.
2. With the **General** tab selected, locate the **Metadata** area.

Metadata support

Nexus Dashboard uses metadata bundles to detect new bugs, PSIRTs, Field Notices, and End of Life Notices. Metadata packages are constantly updated by us and posted to the Cisco Intersight Cloud after validation. Nexus Dashboard connects to the Cisco Intersight Cloud through a device connector that is embedded in the Nexus Dashboard platform and that pulls periodically updated metadata packages. With metadata support for air-gapped environment, if Nexus Dashboard is not connected to Cisco Intersight Cloud, you can manually upload the latest metadata to Nexus Dashboard in a secure and trusted way. You can download the bundle updates from the Cisco DC App Center.

Navigate to **Admin** > **System Settings** > **Metadata** to view the metadata version.

- In the General area, the Metadata Version is displayed.
- In the Update Metadata Version area, you can upload metadata for air-gapped environments.

Metadata support for air-gapped environment

With metadata support for air-gapped environment, if Nexus Dashboard is not connected to Cisco secure cloud, you can upload the latest metadata to Nexus Dashboard periodically in a secure and trusted way.

You can download the encrypted metadata file from the Cisco DC App Center and upload it to Nexus Dashboard to get decrypted updates on exposure to Bugs, PSIRTs, Defects, Field Notices, and End of Life Notices.

Update metadata version

Follow these steps to update the latest metadata version in an air-gapped or offline environment.

1. Log in to Cisco DC App Center.
2. From the User drop-down list, choose **My Account**.
3. Click **Config Files Requests** tab.
4. Click **Request Config File**.
5. From the **Choose App ID** drop-down list, select Nexus Dashboard.

   [request config] | *request-config.jpg*

6. Verify the minimum supported app version and click **Request**.

   It takes approximately 15 minutes for the request to be completed. In the Config Files Request page, the generated file is displayed in the table below.

7. Select the file and click **Download** to download the file locally.

   [request config download] | *request-config-download.jpg*

8. Log in to Nexus Dashboard.
9. Navigate to **Admin** > **System Settings** > **Metadata** to view the metadata version.
10. In the Update Metadata Version area, upload the file you have downloaded from the Cisco DC App Center.
11. Click **Begin Upload** to upload the latest metadata.

Message bus configuration

You might configure a message bus for various reasons, such as when you add a Kafka broker configuration. For more information, see the section "Add Kafka broker configuration" in the Editing Fabric Settings article for your fabric type (for example, Editing Data Center VXLAN Fabric Settings.

To configure the message bus at the **System Settings** level.

1. Navigate to **Admin** > **System Settings** > **General**.
2. In the **Message bus configuration** area, click **Edit**.

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

3. Click **Add message bus configuration**.

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

4. In the **Name** field, enter a name for the configuration.
5. 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.
6. In the **Topic name** field, enter the name of the Kafka topic to which Nexus Dashboard must send the messages.
7. 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.

8. Click **Save**.

Email

Follow these steps to configure email at the system settings level.

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

   The **Email configuration** page appears.

3. Click **Add email configuration**.

   The **Add email configuration** page appears.

4. Enter the necessary information to add the email configuration.

Field	Description
Name	Enter the name for the email configuration.
SMTP host	Enter the SMTP host information. This setting is used as an email out-of-band notification for programmable reports and alarms. The SMTP host address must be reachable through the Nexus Dashboard management interface. If the Nexus Dashboard management interface and SMTP host are part of different IP subnets, then you must create a static route entry in the Nexus Dashboard Cluster configuration.
SMTP host port	Enter the SMTP host port.
Enable SMTP authentication	Check the box to enable SMTP authentication. The default is enabled (box is checked). Enable this option only if your SMTP server requires credential-based authentication. If SMTP authentication is enabled for an SMTP server that does not require or support SMTP AUTH, email delivery can fail.
SMTP authentication username	Enter the SMTP authentication username.
SMTP authentication password	Enter the SMTP authentication password.
Source email	Enter the **From Email** address for reports.

5. Click **Save** in the **Add email configuration** page.

   You are returned to the **Email configuration** page, with your new email configuration listed in the table.

6. Click Save in the **Email configuration** page.

   You are returned to the **System Settings** page.

After you add email configurations at the system settings level, those email configurations become available as a drop-down list in the appropriate pages.

Fabric management

The **Fabric management** page under **System Settings** has these areas:

- Management
- Topology snapshots
- Switch bootstrap
- Advanced settings
- Change control
- Event analytics

Management

Use the **Management** page to make these configurations:

- AAA passthrough of device credentials for LAN fabrics
- ACI Proxy Mode

AAA passthrough of device credentials for LAN fabrics

The authentication, authorization, and accounting (AAA) passthrough feature allows authentication for remote users and eliminates the need to update device credentials manually. This feature enhances the workflow for managing fabrics and switches when a password is changed or updated on the remote server.

The AAA passthrough feature is disabled by default.

Before enabling this feature, ensure that Nexus Dashboard and the switches are configured with the same remote authentication server.

Follow these steps to enable AAA passthrough of device credentials for LAN fabrics.

1. Navigate to **Admin > System Settings > Fabric management > Management**.
2. In the **Management** area, click **Edit**.
3. Check the **Enable AAA passthrough of device credentials** check box.
4. Click **Save**.

ACI Proxy Mode

Follow these steps to enable the ACI proxy mode.

1. Navigate to **Admin > System Settings > Fabric Management > Management**.
2. In the **ACI proxy mode** area, choose the appropriate option:
    - Proxy with ND Service User
    - Disable
    - Proxy with remote users credential

Topology snapshots

Use the **Topology snapshots** area to enable topology snapshots.

1. Navigate to System Settings.
2. Click **Fabric management**.
3. In the **Topology snapshots** area, click **Edit**.
4. Check the box in the **Enable topology snapshots** field to enable this feature.
5. Click **Save**.

Switch bootstrap

You might make configurations in the **Switch bootstrap** area as a prerequisites for secure POAP. Secure POAP is supported from Cisco NX-OS 9000 Release 10.2.3 or higher version switches. For more information, see Configuring Inband Management and Out-of-Band PnP.

1. Navigate to System Settings.
2. Click **Fabric management**.
3. In the **Switch bootstrap** area, click **Edit**.
4. Choose **http**, **https**, or **tftp** from the drop-down list for **Bootstrap script download protocol** field.
    - For the **http** option, you must enter the IP address of the bench router (BR), port number, and name for certificate bundle in **Bench Router URL with port and certificate file name** field. Ensure that the certificates are uploaded on Nexus Dashboard server for values to autopopulate in this field.
    - By default, for the **http** option, the **Bench router URL with port** field in the **Switch bootstrap** window will be blank. After you install the Root CA Certificate bundle on Bench routers, this field will be autopopulated.

      If these fields are autopopulated, with default port number 29151 and URL *https://10.10.10.1:29151/PoapCACertBundle.pem*, you must configure this URL before you install the BR with the Root CA certificate bundle.

    - Make sure that the fabric is in managed mode before configuring the BR.
    - Ensure that you configure the DHCP option if the DHCP server is used.
    - You must upload CA signed POAP server certificate on Nexus Dashboard and upload the corresponding CA certificate bundle for the BR. On Nexus Dashboard, navigate to **Admin > Certificate Management > Fabric certificates** to upload the relevant certificates.

You might also make configurations in the **Switch bootstrap** area when configuring an automatic import of a pre-provisioned device. For more information, see the section "Pre-provision a device" in Working with Inventory in Your Nexus Dashboard LAN and IPFM Fabrics.

1. Navigate to the **Admin > System Settings > Fabric management > Switch bootstrap** dashlet.
2. Check the **Auto admit pre-provisioned switches during re-poap** check box.

By default, the **Auto admin pre-provisioned switches during re-poap** option is

not enabled.

3. Click **Save**.

Advanced settings

The **Advanced** settings area under **System Settings > Fabric management** provides multiple advanced configurations that sometimes differ between fabric types. In addition, you might be directed to the **Advanced settings** area as part of a procedure, where you might have to make a system-level configuration before fabric-level configurations can take effect.

Discovery settings for SAN fabrics

The **Discovery** tab under **Advanced settings** provides system-level configurations for SAN fabric discovery. This section explains AAA remote authentication passthrough and the process for adding organizationally unique IDs (OUIs) to facilitate SAN device discovery.

AAA remote authentication passthrough for SAN fabrics

The authentication, authorization, and accounting (AAA) remote authentication passthrough feature allows authentication for remote users and eliminates the need to update user credentials manually.

When the AAA remote authentication passthrough feature is enabled, it automatically saves discovery and device credentials for remote users after successful login. This feature enhances the workflow for managing fabrics and switches, as users no longer need to change credentials every time a password is changed or updated in the remote server.

If AAA remote authentication passthrough is not enabled and you configure a SAN fabric to use an external credential store, such as CyberArk, Nexus Dashboard retrieves the fabric discovery credentials and device credentials from the credential store by using the configured store key. When AAA remote authentication passthrough is enabled, the AAA remote username and password take precedence and overwrite the credential store key.

The AAA remote passthrough feature is disabled by default.

Before enabling this feature, ensure that Nexus Dashboard and the switches are configured with the same remote authentication server. You must also enable **Display advanced settings and options for TAC support**. For instructions, see Display advanced settings and options for TAC support.

Follow these steps to enable AAA remote authentication passthrough for SAN fabrics.

1. Navigate to **Admin > System Settings > Fabric management > Advanced settings > Discovery**.
2. Check the **Enable AAA Passthrough feature** check box.
3. From the **AAA Passthrough Authentication / Privacy** drop-down list, choose the required option.
4. Click **Save**.

Add OUIs for SAN device discovery

To discover switches with new organizationally unique IDs (OUIs) that might not be detected during initial SAN discovery, enter the switch OUI followed by Cisco m in the **Add New OUIs for Discovering Switch, End Device with Format** field. For example, for a switch with the OUI 0x70a983, enter

0x70a983 Cisco m, and then click **Save**.

Change control

For more information on change control, see Using Change Control and Rollback in Your Nexus Dashboard.

The change control feature is disabled by default.

Follow these steps to enable the change control feature.

1. Navigate to the **Fabric Management** page in Nexus Dashboard:

   Admin > System Settings > Fabric Management

2. Locate the **Change control** feature box in the **Fabric Management** page and click **Edit** in that box.
3. Click the box next to **Enable change control** to enable the change control feature.
4. Determine if there are other configuration settings that you want for change control.
   - **Enable for Orchestration**: Check the box to enable change control for orchestration for ACI fabrics. The orchestration options are not supported for NX-OS fabrics.

     The following options become available if you check the **Enable for Orchestration** option.

     - **Required number of approvers**: Enter the number of approvers required for change control tickets. Default is 1 approver.
     - **Allow self approval**: This field is enabled by default. Check the box to allow users who select or create a change control ticket for an action to also approve that change control ticket.
   - **Enable for ND Managed Fabrics**: Check the box to enable change control for Nexus Dashboard-managed fabrics.

     The following options become available if you check the **Enable for ND Managed Fabrics** option.

     - **Enable bypass change control for telemetry**: Check the box to enable bypass change control for telemetry.

       If there are any system-triggered changes to configuration settings related to telemetry, a ticket approval or deployment process is not required, as these changes will be automatically approved and deployed if the **Enable bypass change control for telemetry** option is enabled.

       

       You can also disable **Enable bypass change control for telemetry** option, if you want to create a ticket to approve/deploy. The disable option will not work for colocation fabrics.

     - **Ticket name prefix**: Change the Ticket ID prefix string from the default TICKET_, if necessary. See Configuring the Auto Generated Ticket ID for more information.

5. Click **Save**.

Once you have enabled the change control feature, all areas that are supported with change control are now tracked with a ticket. See Guidelines and Limitations: Change Control for a list of areas that are tracked with a ticket when change control is enabled, and see Typical Change Control Workflow to understand how having the change control feature enabled affects the Nexus Dashboard operations that are supported with change control.

You can return to the **Fabric Management** page to disable the change control feature; however, disabling the change control feature after you have enabled it is not supported when active tickets (works in progress) are detected. You must complete the change control process on any active tickets before you can disable the change control feature in this case.

Once you have successfully disabled the change control feature, the change operations will no longer be tracked with change control tickets. Completed tickets and associated data remains, but you cannot perform any actions against those areas.

Event analytics

Follow these steps to enable the event analytics feature.

1. Navigate to the **Fabric Management** page in Nexus Dashboard:

   Admin > System Settings > Fabric Management

2. Locate the **Event analytics** feature box in the **Fabric Management** page and click **Edit** in that box.

Events

This tab displays the events that are generated for the switches. You can select one or more events and then acknowledge or unacknowledge their status using the Change Status drop-down list. In addition, you can select one or more alarms and then click the Delete button to delete them. If you want to delete all events, click the Delete All button.

The following table describes the fields that appear on **Analyze > Event Analytics > Events**.

Field	Description
Group	Specifies the Fabric
Switch	Specifies the hostname of the switch
Severity	Specifies the severity of the event
Facility	Specifies the process that creates the events. The event facility includes two categories: Nexus Dashboard and syslog facility. Nexus Dashboard facility represents events generated by Nexus Dashboard internal services and SNMP traps generated by switches. Syslog facility represents the machine process that created the syslog messages.
Type	Specifies how the switch/fabric are managed
Count	Specifies the number of times the event has occurred

Field	Description
Creation Time	Specifies the time when the event was created
Last Seen	Specifies the time when the event was run last
Description	Specifies the description provided for the event

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

Action Item	Description
Acknowledge	Select one or more events from the table and choose **Acknowledge** icon to acknowledge the event information for the fabric. After you acknowledge the event for a fabric, the acknowledge icon is displayed in the Ack column next to the Group.
Unacknowledge	Select one or more events from the table and choose **Unacknowledge** icon to acknowledge the event information for the fabric.
Delete	Select an event and choose *Delete*to delete the event.
Add Suppressor	Select an event and choose **Add Suppressor** to add a rule to the event. You can provide name to the rule. Using the **Scope** options, you can add this rule to all the Fabrics, or particular elements or all elements.

Setting Up Events

Follow these steps to setup an event using the Cisco Nexus Dashboard Web UI.

1. Click **Receiver**.

   The **Receiver** tab displays the following details:

   - **Syslog Receiver enabled**: Displays the status of the syslog server.
   - **SNMP Trap Receiver**: Displays the details of SNMP traps received, processed and dropped.
   - **Syslog Receiver**: Displays the details of syslog messages received, processed and dropped.

- Ensure that you allow access to the SMTP service from the cisco-ndfc-dcnm-syslog-trap interface for forwarding of event email notifications. The cisco-ndfc-dcnm-syslog-trap interface also provides access to the switches for SNMP queries. For more information, see the section "Nexus Dashboard persistent IP addresses" in Cisco Nexus Dashboard Deployment and Upgrade Guide.
- The cisco-ndfc-dcnm-syslog-trap interface also provides access to the switches for SNMP. Ensure that you allow access to the SNMP destination port 161 on the switches for the cisco-ndfc-dcnm-syslog-trap interface. For more information, see the section “Communication Ports for LAN deployments in Cisco Nexus Dashboard Deployment and Upgrade Guide.

2. Click the **Sources** tab to view a list of fabrics and their associated switches.

   The **Sources** tab displays all the fabrics and the associated switches in tabular format. It also displays if traps and syslogs have been configured on the switches.

3. Perform the following steps to create rules for forwarding email notifications or traps for events:

   Cisco Nexus Dashboard Web UI forwards fabric events through email or SNMPv1 or SNMPv2c traps. Some SMTP servers require SMTP authentication parameters, while others do not support or require SMTP AUTH. Enable SMTP authentication in Nexus Dashboard only when the target SMTP server requires it. If SMTP authentication is enabled for a server that does not require or support it, event email notifications can fail to be delivered.

   a. Ensure that you have configured SMTP parameters before configuring rules for forwarding event notifications through emails.

      To verify SMTP configuration, navigate to **Admin > System Settings > General > Email** and verify that you have configured the required fields.

   b. To configure rules, click the **Forwarding** tab and choose **Actions > Add Rule** and configure the fields as described in the following table.

      Configure Rules

Field	Description
Forwarding Method	Choose one of the forwarding methods: • **E-Mail** • **Trap**
Email Address	This field appears if you select **E-mail** as the forwarding method. Enter an email address for forwarding the event notifications.
Fabric	Select **All LAN Groups** or a specific fabric for notification.
Source	Select **NDFC** or **Syslog**. If you select **DCNM**, do the following: 1. From the **Type** drop-down list, choose an event type. 2. Check the **Storage Ports Only** check box to select only the storage ports. This check box is enabled only for port related events. If you select **Syslog**, do the following: 1. In the **Facility** list, select the syslog facility. 2. In the **Type** field, enter the syslog type. 3. In the **Description Regex** field, enter a description that matches with the event description.

   c. From the **Minimum Severity** drop-down list, select the severity level of the messages to receive.

The traps that are transmitted by Cisco Nexus Dashboard correspond to the severity type. A text description is also provided with the severity type.

```
trap type(s) = 40990 (emergency)
40991 (alert)
40992 (critical)
40993 (error)
40994 (warning)
40995 (notice)
40996 (info)
40997 (debug)
textDescriptionOid = 1, 3, 6, 1, 4, 1, 9, 9, 40999, 1, 1, 3, 0
```

d. Click **Add Rule**.

4. Follow these steps to create rules for suppressing events.

Nexus Dashboard allows you to suppress specified events based on user-specified rules. Such events will not be displayed on the Nexus Dashboard Web UI and client. The events will neither be added to the Nexus Dashboard database, nor forwarded via email or as SNMP traps.

You can view, add, modify, and delete rules from the table. You can create a rule from the existing events. Select an existing event as the template and open the **Add Rule** window by navigating to **Events** page, select the event and choose **Actions > Add Suppressor**. The details are automatically ported from the selected event in the events table to the fields of the **Add Rule** window.

a. In the **Name** field, enter a name for the rule.

b. In the **Scope** field, select one of the following options - **SAN**, **Port Groups** or **Any**.

In the **Scope** field, the LAN/SAN groups and the port groups are listed separately. For SAN and LAN, select the scope of the event at the fabric or group or switch level. You can only select groups for port group scope. If use select **Any** as the scope, the suppression rule is applied globally.

c. In the **Facility** field, enter the name or choose from the SAN/LAN switch event facility list.

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

d. In the **Type** field, enter the event type.

If you do not specify the event type, wildcard is applied.

e. In the **Description Matching** field, specify a matching string or regular expression.

The rule matching engine uses regular expression that is supported by Java Pattern class to find a match against an event description text.

f. Check the **Active Between** check box and select a valid time range during which the event is suppressed.

By default, the time range is not enabled.

In general, you must not suppress accounting events. Suppression rule for Accounting events can be created only for certain situations where accounting events are generated by actions of Nexus Dashboard or switch software. For example, 'sync-snmp-password' AAA syslog events are automatically generated during the password synchronization between Nexus Dashboard and managed switches. To suppress accounting events, navigate to the **Events** page, select the event and choose **Actions > Add Suppressor**.

g. Click **Add Rule**.

Multi-cluster connectivity

Use the **Multi-cluster connectivity** page to connect to another Nexus Dashboard cluster for a single pane of glass view into all cluster's fabrics and services.

You might use multi-cluster connectivity in several situations, such as:

- Connecting Nexus Dashboard clusters. See Connecting Nexus Dashboard clusters for more information.
- Connecting APIC clusters. See Connecting ACI clusters for more information.
- Configuring a primary cluster as the default authentication domain. See the section "Configure primary cluster as the default authentication domain" in Configuring Users, Roles, and Security for more information.

Connecting Nexus Dashboard clusters

These sections provide the necessary information to connect Nexus Dashboard clusters:

- Guidelines and limitations: Nexus Dashboard cluster connectivity
- Connect multiple Nexus Dashboard clusters
- Disconnect Nexus Dashboard clusters

Guidelines and limitations: Nexus Dashboard cluster connectivity

The following guidelines apply when configuring Nexus Dashboard multi-cluster connectivity:

- Only users with a 'super-admin' role and all security domain can add or delete a Nexus Dashboard cluster. See Configuring Users and Security for more information.
- When configuring multi-cluster connectivity in Nexus Dashboard 4.1.1, you can only connect clusters running on Nexus Dashboard 4.1.1 and later. If any cluster that is part of the multi-cluster connectivity configuration is running on Nexus Dashboard 4.1.1, the other clusters in that multi-cluster connectivity configuration must also be running on Nexus Dashboard 4.1.1 or later.
- For supported scale limits, such as number of clusters that can be connected together and number of fabrics across all clusters, see the *Nexus Dashboard Release Notes* for your release.
- Connectivity (HTTPS) must be established between the management interfaces of all the nodes of all the clusters, which will be connected via multi-cluster connectivity.
- For multi-cluster fabric group operations, the following table lists the supported and unsupported IP stack combinations.

Cluster IP stack combination	Supported
Single IPv4 stack with dual-stack	Supported
Single IPv6 stack with dual-stack	Supported
Single IPv4 stack, single IPv6 stack, and dual-stack	Not supported

- The names of the fabrics onboarded in the clusters that you plan to connect together must be unique across those clusters.

  Duplicate fabric names across different clusters may result in DNS resolution failures.

- The primary cluster, which you use to establish multi-cluster connectivity, must be running the same or a later release of Nexus Dashboard than any other cluster in the group.

  In other words, you cannot connect a Nexus Dashboard cluster running release 2.3.1 from a primary cluster that is running release 3.0.1.

- If you are upgrading multiple clusters that are connected together, you must upgrade the primary cluster first.
- From any cluster in the connected clusters group, you can view other clusters only if they are running the same or earlier version of Nexus Dashboard.

  In other words, if cluster1 is running release 2.3.1 and cluster2 is running release 2.2.1, you can view cluster2 from cluster1 but not vice versa.

- You can either use the primary cluster as an authentication domain, see the "Multi-cluster primary as an authentication domain" section in Configuring Users and Security, or you can use a remote user.

  If you connect multiple clusters but login to one of the clusters as a local admin user, you will only be able to view and manage the local cluster into which you logged in.

  To view and manage all clusters in the group, you must login as a remote user configured on all clusters.

- In a multi-cluster environment, ACI fabrics can be onboarded from the cluster that has connectivity to the fabric, including a secondary cluster.

  For ACI-specific onboarding behavior, fabric visibility, and deregistration details, see Connecting ACI clusters.

Guidelines and limitations: Cluster names when restoring a backup

In releases prior to Nexus Dashboard release 4.3.1, certain isolated restrictions might cause issues in situations where they come into play together. For example, these two isolated restrictions:

- If a backup was taken from a federated cluster, then the cluster where you were restoring had to have the same cluster name; otherwise, the restore would fail.
- You cannot change a cluster name on a running cluster.

cause issues together when you want to set up a backup of a cluster in a federation.

As an example, assume that you have clusters alpha, beta, and gamma, where

- clusters alpha and beta are in a federation, where alpha is the primary cluster and beta is the secondary cluster, and
- cluster gamma is a standalone cluster.

Now, assume that you backed up the beta cluster (the secondary cluster in the federation), and the

beta cluster fails. You could then restore the beta cluster on the standalone gamma cluster. After a successful restore, the gamma cluster is now restored with beta cluster settings.

Ideally, you would want to bring the gamma cluster back into the federation as the beta cluster because, as part of a federated cluster, the backup cluster and the cluster where you are restoring have to have the same cluster name; however, you can't rename the gamma cluster to beta because you can't change the name of a running cluster.

Beginning with Nexus Dashboard release 4.3.1, for both federated clusters and standalone clusters, when restoring a backup, the local cluster name is now changed to match the backup.

As an example of the new workflow, using the same alpha, beta, and gamma cluster example from above, assume that you backed up the beta cluster (the secondary cluster in the federation), and the beta cluster fails.

- You can now restore the beta cluster on the standalone gamma cluster,
- gamma is renamed to beta during the restore process, and
- you can re-register the beta cluster into the federation.

As another example, assume that you backed up the alpha cluster (the primary cluster in the federation), and the alpha cluster fails.

- You could then restore the alpha cluster on the standalone gamma cluster,
- gamma is renamed to alpha during the restore process, and
- you can re-register the beta cluster from the new alpha cluster.

Connect multiple Nexus Dashboard clusters

Before you begin

- Familiarize yourself with the information provided in the Guidelines and limitations: Nexus Dashboard cluster connectivity section.
- Set up remote authentication and users on all clusters that you plan to connect.

  Multi-cluster connectivity is supported for remote users only, so you must configure the same remote user with admin privileges for all clusters. For additional details, see the "Remote authentication" section in Configuring Users and Security.

To connect another cluster:

1. Log in to the Nexus Dashboard GUI of the cluster which you want to designate as the primary.
2. Add the second cluster.
   a. From the main navigation menu, choose **Admin > System Settings**.
   b. In the main pane, click **Multi-cluster connectivity**.
   c. Click **Connect Cluster**.
   d. In **Select type**, choose **Nexus Dashboard**.
   e. Click **Next**.

      You advance to the **Settings** step in the **Connect Cluster** workflow.

3. Provide the necessary cluster information.
    a. In the information fields, provide the hostname or IP address and the authentication information for the cluster you are adding.

       You only need to provide the management IP address of one of the nodes in the target cluster. Other nodes' information will be automatically synced after connectivity is established.

       - The user you provide must have administrative rights on the cluster you are adding. The user credentials are used once when you are first establishing connectivity to the additional cluster. After initial connectivity is established, all subsequent communication is done through secure keys. The secure keys are provisioned to each cluster while adding it to the group.
       - The cluster you are adding must not be part of an already existing group of clusters.

    b. When you have entered all of the necessary configuration information, click **Next**.

       You advance to the **Summary** step in the **Connect Cluster** workflow.

4. Verify all of the information that is shown in the summary page is correct.
5. If all of the information shown in the page looks correct, click **Submit**.
6. Repeat the procedure for any additional Nexus Dashboard cluster which you want to add to the group.

   After multiple clusters are added to the group, you can see their status in the **Cluster Configuration > Multi-cluster connectivity** page.

   Note that while you can view and manage any cluster from any other cluster as long as they are part of the same multi-cluster group, you can only add and remove clusters from the group when viewing the primary cluster.

The **Multi-cluster connectivity** page displays all clusters that are part of the multi-cluster group. The **Connect Cluster** button is shown only when viewing the primary cluster. To modify the cluster group, you need to navigate to the primary cluster, at which point the **Connect Cluster** button becomes available:

- The **Cluster: <name>** dropdown in the main navigation menu shows the cluster you are currently viewing.

  You can select a different cluster from this dropdown, which opens a new page allowing you to navigate to another cluster in the same group.

  

  While the 2.x releases of Nexus Dashboard allowed you to view and manage any cluster from any other cluster as long as they were part of the same multi-cluster group, release 3.0.1 changed this behavior. You can now easily navigate between clusters by picking a specific cluster from the **Cluster** dropdown in the main navigation pane, but you cannot manage or configure another cluster directly from the one where you are logged in.

  - The Primary label indicates the group's primary cluster.

    You must be viewing this cluster to make any changes to the cluster group, such as adding or removing clusters.

- The Local label indicates the cluster you logged into.

  This is the cluster whose address is displayed in the browser's URL field. If you navigate to a different cluster as mentioned above, the browser URL and the Local label will not change.

- The Connectivity Status: Shows the status of the uplink to the cluster.
- The URL shows the list of IP addresses of the cluster.
- The **Actions** (...) menu for each cluster allows you to **Re-Register** and **Disconnect Cluster**

- The **Connect Cluster** allows you to add a new cluster.

Disconnect Nexus Dashboard clusters

To disconnect a cluster from an existing group:

1. Log in to the Nexus Dashboard GUI of the primary cluster.

   Adding and removing clusters from the group must be done from the primary cluster.

2. From the main navigation menu, select **Admin > System Settings**.
3. In the main pane, select **Multi-cluster connectivity**.
4. From the **Actions** (...) menu for the cluster you want to remove, select **Disconnect Cluster**.
5. If the cluster status is still shown as Up at this time, you will be given an option to forcefully remove the member. This option should be used only if previous removal attempts were unsuccessful.
6. In the confirmation page, click **Ok**.

- When the primary cluster becomes unreachable from the secondary cluster, you can remove the secondary cluster either through the Nexus Dashboard GUI or the API:

  To remove the secondary cluster through the Nexus Dashboard GUI:

  a. Log in to the Nexus Dashboard GUI of the secondary cluster.
  b. Navigate to **Admin > System Settings > Multi-cluster connectivity**.
  c. Select the **Local** cluster.
  d. Click **Disconnect from multi-cluster**.

  To remove the secondary cluster through the Nexus Dashboard API, use the external API /api/v1/infra/clusterActions/disconnectFromMultiCluster. See the API Reference for the full payload and HTTP method.

- To add a primary cluster that has no secondary clusters to a different multi-cluster group, you need to first reset the multi-cluster state from the primary cluster. You can do this either through the Nexus Dashboard GUI or the API:

  To remove the multi-cluster state from the primary cluster through the Nexus Dashboard GUI:

  a. Log in to the Nexus Dashboard GUI of the primary cluster.

b. Navigate to **Admin > System Settings > Multi-cluster connectivity**.

c. Select the **Local** cluster.

d. Click **Disconnect from multi-cluster**.

To remove the multi-cluster state from the primary cluster through the Nexus Dashboard API, use the external API /api/v1/infra/clusterActions/disconnectFromMultiCluster. See the API Reference for the full payload and HTTP method.

Connecting ACI clusters

These sections provide the necessary information to connect Cisco Application Centric Infrastructure (ACI) clusters:

- Guidelines and limitations: ACI cluster connectivity
- Connect ACI clusters
- Disconnect ACI clusters

Guidelines and limitations: ACI cluster connectivity

- Even though you can have multiple Nexus Dashboard clusters with the same name, you cannot have a single ACI fabric in multiple Nexus Dashboard clusters if those clusters have the same name. For example, if you have two Nexus Dashboard clusters where both Nexus Dashboard clusters are named nexus, you cannot add the same ACI fabric to both of those nexus Nexus Dashboard clusters.
- Only users with a Fabric Administrator role and the all security domain access can add and delete ACI clusters. See Configuring Users and Security for more information.
- You can onboard an ACI fabric from either the primary cluster or a secondary cluster.
- If more than one cluster in the multi-cluster group has connectivity to the ACI fabric, onboard the fabric from the cluster on which you want to enable and manage the fabric services.
- After any cluster in the multi-cluster group onboards the ACI fabric, the fabric becomes available in the multi-cluster environment.
- The primary cluster can update the License tier, Security domain, and Location settings for the fabric even when the fabric is not reachable.
- When you onboard an ACI cluster to Nexus Dashboard, the APIC might detect previous Nexus Dashboard cluster registrations that conflict with your current local Nexus Dashboard cluster. This conflict can happen if a previous Nexus Dashboard cluster was registered with a node serial number or in-band IP address that matches one from your current local Nexus Dashboard cluster.

  If you are unable to onboard an ACI cluster due to this issue, the system will return the names of all conflicting Nexus Dashboard clusters. This issue should resolve automatically when you onboard ACI through the Nexus Dashboard GUI using the procedures in Connect ACI clusters, where Nexus Dashboard will help to clean up conflicting ACI registrations. However, if the Nexus Dashboard GUI access is unavailable, use the following procedures to manually clean up conflicting APIC registrations.

  For each cluster name returned, take action based on the applicable scenario:

- The conflicting cluster is an active Nexus Dashboard cluster with an IP overlap with the current local Nexus Dashboard cluster, do one of the following:
  - if that Nexus Dashboard cluster does not need access to the ACI cluster, you can unregister the ACI cluster with that Nexus Dashboard cluster.

    or

  - Redeploy the current Nexus Dashboard cluster with different in-band IP addresses.
- The conflicting cluster is permanently inactive. This can happen if a cluster is re-deployed with a new cluster name without deleting its ACI fabrics first.
  - The conflicting cluster registration must be deleted from the ACI cluster using the following REST calls:

    ```
    POST https://<APIC>/api/aaaLogin.json
    payload:
    {
       " aaaUser" : {
          " attributes" : {
             " name" : " admin" ,
             " pwd" : " myPassword"
          }
       }
    }
    ```

    The response body will contain a token. Pass this token as the Cookie header in the next request, using the key Cookie and value APIC-Cookie=token.

  - To delete each cluster:

    ```
    DELETE https://<APIC>/api/mo/uni/userext/snclstr-
    <ND_CLUSTER_NAME_TO_BE_REMOVED>.json
    ```

- Fabric connectivity must be already configured as described in the section " Fabric Connectivity" in the *Cisco Nexus Dashboard and Services Deployment and Upgrade Guide*.
- EPG/L3Out for Nexus Dashboard data network IP connectivity must be already configured as described in the section " Fabric Connectivity" in the *Cisco Nexus Dashboard and Services Deployment and Upgrade Guide*.
- IP connectivity from Nexus Dashboard to ACI cluster in-band IP over the data network must be already configured.
- IP connectivity from Nexus Dashboard to the leaf nodes' and spine nodes' in-band IPs over the data network must be already configured.

Connect ACI clusters

To onboard one or more Cisco ACI fabrics or clusters to your Nexus Dashboard in a multi-cluster

environment:

1. Log in to the Nexus Dashboard GUI of the cluster from which you want to onboard the ACI fabric.

   

   In a multi-cluster environment, you can onboard the ACI fabric from any cluster in the multi-cluster group that has connectivity to the fabric, including a secondary cluster. After onboarding is complete, the fabric becomes available in the multi-cluster environment.

2. Add the ACI fabric or cluster.
   a. From the main navigation menu, select **Admin > System Settings**.
   b. In the main pane, select **Multi-cluster connectivity**.
   c. Click **Actions > Connect cluster**.

      The **Connect Cluster** page appears.

      

      You are also redirected to this **Connect Cluster** page if you are creating an ACI fabric through **Manage > Fabrics > Local > Create Fabric > Onboard ACI Fabric**.

   d. In the **Select type** page, choose **ACI**.
   e. Click **Next**.

      You advance to the **Settings** step in the **Connect Cluster** workflow.

3. Provide the necessary ACI fabric or cluster information.
   - **Host Name/IP Address**—Provide the IP address used to communicate with the Cisco ACI.

     

     When providing the address, do not include the protocol (http:// or https://) as part of the URL string or fabric onboarding will fail.

   - **User Name and Password**—Login credentials for a user with admin privileges on the fabric you are adding.
   - (Optional) **Login Domain**—If you leave this field empty, the fabric's local login is used.
   - (Optional) **Validate peer certificate**—Allows Nexus Dashboard to verify that the certificates of hosts to which it connects (such as fabric controllers) are valid and are signed by a trusted Certificate Authority (CA).

     

     You must have the certificate for this fabric already imported into your Nexus Dashboard before you can add a fabric using this option. If you have not yet added the certificates, cancel the onboarding workflow and follow the instructions described in the "Administrative Tasks" article in the Nexus Dashboard documentation library; then after you have imported the certificates, add the fabric as described here. If you enable the Verify Peer Certificate option but don't import the valid certificate, fabric onboarding will fail.

4. When you have entered all of the necessary configuration information, click **Next**.

   You advance to the **Onboard fabric** step in the **Connect Cluster** workflow.

5. Configure the parameters and capabilities of the ACI fabric.

Field	Description
Fabric Name	Enter a unique name for the fabric.
Location	Choose the location for the fabric.
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.
Enable telemetry	Check the box to enable **Telemetry** for the fabric. This is the equivalent of enabling the Nexus Dashboard Insights service in previous releases.
Telemetry collection	This option becomes available if you choose to enable **Telemetry** in the **Enabled features** field above. Choose either **Out-of-band** or **In-band** for telemetry collection. Regardless of the option that you choose for telemetry collection, reachability to the ACI IP address must be through the Nexus Dashboard data interface.
Telemetry streaming	This option becomes available if you choose to enable **Telemetry** in the **Enabled features** field above. Choose either **IPv4** or **IPv6** for telemetry streaming.
Security domain	Choose the security domain for the fabric.

6. When you have entered all of the necessary configuration information, click **Next**.

   You advance to the **Summary** step in the **Connect Cluster** workflow.

7. Verify all of the information that is shown in the summary page is correct.
8. If all of the information shown in the page looks correct, click **Connect**.
9. Verify that the onboarded ACI fabric appears in the multi-cluster environment and is available to the connected clusters.

   The location in which the fabric is displayed in the Nexus Dashboard GUI can differ depending on cluster role and fabric state.

10. Repeat the procedure for any additional ACI fabric or cluster which you want to add to the group.

After you have added all of the ACI fabrics or clusters to the group, you can see their status in the **Admin > System Settings > Multi-cluster connectivity** page. You can also cross-launch these onboarded ACI clusters from the Nexus Dashboard GUI by clicking on the onboarded ACI fabric through **Manage > Fabrics** in the Nexus Dashboard GUI, and then clicking on **Open fabric** in that ACI fabrics **Overview** page.

Beginning with Nexus Dashboard release 4.1.1, when you are in an ACI cluster that has been onboarded to Nexus Dashboard, if that ACI is running on release 6.1.4 or later, you can also cross-launch from that ACI's GUI back to the Nexus Dashboard where it's onboarded. For more information, see Nexus Dashboard cluster from APIC GUI.

Re-register clusters

If you upgrade your Nexus Dashboard from an earlier release to Nexus Dashboard release 4.1.1, after the Nexus Dashboard upgrade to release 4.1.1 is complete, you will have to re-register the onboarded ACIs to use the ACI-to-Nexus Dashboard cross-launch functionality introduced in Nexus Dashboard release 4.1.1.

To re-register clusters:

1. Navigate to **Admin > System Settings > Multi-cluster connectivity**.
2. Choose the appropriate ACI cluster, then click **Actions > Re-register cluster**.

Disconnect ACI clusters

To disconnect a cluster from an existing group:

1. If this is a multi-cluster environment and any secondary cluster has registered the ACI fabric or has services enabled for the fabric, first log in to each affected secondary cluster, disable any services enabled for that fabric, and use Deregister to clean up the registration between that Nexus Dashboard cluster and the ACI fabric. After all required secondary-cluster deregistrations are complete, remove the ACI fabric from the primary cluster.
2. Log in to the Nexus Dashboard GUI of the primary cluster.

   Adding and removing clusters from the group must be done from the primary cluster.
3. From the main navigation menu, select **Admin > System Settings**.
4. In the main pane, select **Multi-cluster connectivity**.
5. From the **Actions** (...) menu for the cluster you want to remove, select **Disconnect Cluster**
6. If the cluster status is still shown as Up at this time, you will be given an option to forcefully remove the member. This option should be used only if previous removal attempts were unsuccessful.
7. In the confirmation page, click **Ok**.

- You can disconnect nodes from the multi-cluster group only through the main cluster. If the primary cluster is unavailable, some multi-cluster management operations can still be restricted. However, secondary-cluster onboarding and secondary-cluster deregistration behavior for ACI fabrics are supported as described in this article. To remove a secondary cluster from the group when the primary cluster is unavailable, you must use the

/api/v1/infra/clusters/<primary_cluster_name>/remove API call on the secondary cluster. For more information, see the API Reference.

- In previous releases, when you disconnected all member clusters, the primary cluster would also remove itself from a multi-cluster group. Beginning with Nexus Dashboard release 4.1.1, you must now manually remove the primary cluster from a multi-cluster group.

Deregister a secondary cluster from an ACI fabric

An ACI fabric registration with a Nexus Dashboard cluster is unnecessary when the cluster has no features enabled on the fabric. The Deregister action can be used to clean up the registration between one Nexus Dashboard cluster and an ACI fabric.

Deregistering a fabric on a secondary cluster in a multi-cluster environment removes the fabric from the cluster's **Multi-Cluster connectivity** page. After deregistering a fabric, the registration can be restored by navigating to the **Remote fabrics** tab on the main **Fabrics** page.

Before you begin

- Verify that no services are enabled for the ACI fabric on the secondary cluster.

Procedure

1. Log in to the Nexus Dashboard GUI of the secondary cluster.
2. Navigate to **Admin > System Settings > Multi-Cluster connectivity**.
3. Select the ACI fabric that you want to deregister.
4. From the available actions, choose Deregister.
5. When prompted, enter the admin credentials for the fabric, and then confirm the deregistration action.

Flow collection

The **Flow collection** page under **System Settings** has Flow collection mode.

Flow collection mode

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 must manually configure a Flow Telemetry rule to acquire comprehensive end-to-end information about the troublesome flow.

Beginning with Nexus Dashboard 4.2.1, **Flow telemetry** and **Traffic analytics** capabilities are enhanced to provide granular visibility into micro-segmented environments within Cisco NX-OS and ACI fabrics with security groups configured. This feature enriches flow records and service visualizations with microsegmentation endpoint groups (uSeg EPG) information, offering deeper insights into traffic patterns and security policy enforcement.

These cards are available under Flow collection mode:

- Traffic analytics
- Flow telemetry

Limitations for endpoint security group name enrichment

These are the limitations for endpoint security group (ESG) name enrichment:

- Traffic analytics full mode: The system does not display the endpoint group (EPG) name for external endpoints routed through L3Outs.
- Traffic analytics compatibility mode: The system does not display the ESG name for external endpoints.
- Flow telemetry: The system does not display the EPG name for external endpoints routed through L3Outs.
- NetFlow: The system does not display the ESG name for external endpoints.

Traffic analytics

Use traffic analytics to automatically discover services and visualize flows based on Layer 4 ports, identifying congestion, latency, drops and more. Flow collection jobs are not supported on fabrics with out-of-band streaming.

You must disable Flow Telemetry at the fabric level for every fabric before you can modify these settings. Even though Flow Telemetry and Flow Telemetry Events are enabled at the fabric level, the settings are applied cluster-wide. See the editing fabric settings article for your fabric type for those procedures.

To enable traffic analytics:

1. Navigate to System Settings.
2. Click **Flow collection**.
3. Click **Traffic analytics** under **Flow collection mode**.

   Changing the Flow collection mode will erase all flow records previously collected on all fabrics. In the **Change Flow collection** mode popup, click **Change and save**.

For information, see Analyzing and Troubleshooting Your Network.

Flow telemetry

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**.

In NX-OS fabrics, flows for security groups excludes reserved security group tags (SGTs) 0 to 15.

PTP

This section explains the functionality of the Precision Time Protocol (PTP).

Precision Time Protocols (PTP) is an industry-standard protocol used to synchronize clocks of multiple devices in a network with sub-microsecond accuracy. PTP is crucial in applications where time accuracy is critical and is designed to provide accurate time synchronization for devices in both local area networks (LANs) and wide area networks (WANs), including the Internet. PTP is a critical requirement for the telemetry functions of Nexus Dashboard to provide accurate metrics on latency, congestion and performance.

Properly configuring PTP on your system allows you to avoid having inaccurate traffic analytics monitoring, which might provide false alarms and false calculations of latencies. For this reason, you must enable PTP on the fabric before you can enable **Flow telemetry** or **Traffic analytics** under **Admin > System Settings > Flow collection**.

These sections give additional information on PTP:

- Guidelines and limitations: PTP
- Configure PTP
- PTP-related anomalies

Guidelines and limitations: PTP

- PTP is required on any fabric where Nexus Dashboard monitors flow telemetry or traffic analytics.
- Nexus Dashboard can monitor PTP configuration and telemetry for NX-OS fabrics only.
- Nexus Dashboard strictly enforces a check for NX-OS fabrics. In order to enable flow telemetry or traffic analytics, you must enable PTP on the fabric. This check is enforced for the following fabric

types:

- Data Center VXLAN EVPN (with either iBGP or eBGP overlay routing protocol)
- AI Data Center VXLAN
- AI Routed
- Classic LAN
- Routed

- Nexus Dashboard can monitor the status of PTP in NX-OS fabrics. The PTP status is visible in the Nexus Dashboard telemetry status page.
- PTP monitoring through telemetry is supported in NX-OS release 9.3(7) and later.
- Nexus Dashboard cannot monitor the status of PTP in ACI fabrics. Instead, only significant PTP correction events across the ACI fabric are detectable through APIC faults, which are then translated into Nexus Dashboard anomalies. In addition, for ACI fabrics, the PTP status is not visible in the Nexus Dashboard telemetry status page.
- If you have a Nexus Dashboard running on release 3.2.x and you want to upgrade to Nexus Dashboard release 4.1.1:
  - If you have flow collection enabled without PTP also configured on your 3.2.1/3.2.2 system, when you upgrade to Nexus Dashboard release 4.1.1, the PTP data cannot be streamed correctly through telemetry until you properly configure and enable PTP on the physical fabric, either through the controller or manually.
  - If you do not have flow collection enabled on your 3.2.1/3.2.2 system, when you upgrade to Nexus Dashboard release 4.1.1, if you attempt to enable flow collection without also enabling PTP on your 4.1.1 system, you will see an error message and you will not be able to proceed until you properly configure and enable PTP.

Configure PTP

The procedures for configuring PTP differs, depending on the type of fabric being managed by Nexus Dashboard:

- Configure PTP: NX-OS fabrics
- Configure PTP: ACI fabrics

Configure PTP: NX-OS fabrics

To configure PTP in an NX-OS fabric:

1. Enable PTP at the fabric level.
   a. Navigate to the main **Fabrics** window:

      Manage > Fabrics

   b. Locate the fabric that you want to edit.
   c. Click the circle next to the fabric that you want to edit to select that fabric, then click **Actions > Edit fabric settings**.

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

d. Click **Fabric management > Advanced**.

e. Locate the **Enable Precision Time Protocol (PTP)** field and click the checkbox.

This enables PTP across a fabric. When you check this check box, PTP is enabled globally and on core-facing interfaces.

f. Configure the associated PTP fields that became editable when you enabled the **Enable Precision Time Protocol (PTP)** option.

Field	Description
PTP Source Loopback Id	Specifies the loopback interface ID Loopback that is used as the Source IP Address for all PTP packets. The valid values range from 0 to 1023. The PTP loopback ID cannot be the same as RP, Phantom RP, NVE, or MPLS loopback ID. Otherwise, an error will be generated. The PTP loopback ID can be the same as BGP loopback or user-defined loopback which is created from Nexus Dashboard. If the PTP loopback ID is not found during **Deploy Config**, the following error is generated: Loopback interface to use for PTP source IP is not found. Create PTP loopback interface on all the devices to enable PTP feature.
PTP Domain Id	Specifies the PTP domain ID on a single network. The valid values range from 2 to 127.
PTP Source VLAN Id	Specifies the SVI used for PTP source on ToRs. The valid values range from 2 to 3967.

g. Click **Save** to save the PTP configurations.

2. Enable PTP on the core-facing interfaces using the freeform config template.

We also recommend that you add the ttag/ttag-strip entries to prevent any ttag frame drop in the core network. For more information, see the section "Enable freeform configurations on fabric switches" in Working with Inventory in Your Nexus Dashboard LAN or IPFM Fabrics.

Following is an example configuration:

```
feature ptp
ptp source 100.100.100.10 // IP address of the loopback interface (loopback0) that is already created or user created loopback interface in the fabric settings
ptp domain 1 // PTP domain ID specified in fabric settings
interface Ethernet1/59 // Core facing interface ptp
interface Ethernet1/50 // Host facing interface
ttag
ttag-strip
```

3. Enable **Flow telemetry** or **Traffic analytics** at the system level and fabric level as you normally

would.

- At the system level:
  a. Navigate to **Admin > System Settings > Flow collection**.
  b. Enable **Flow telemetry** or **Traffic analytics** at the system level in this page.

     See Analyzing and Troubleshooting Your Network for more information.

- At the fabric level:
  a. Navigate to **Manage > Fabrics**.
  b. Choose the appropriate fabric and click **Actions > Edit Fabric Settings**.
  c. Click **Telemetry**, then enable **Flow telemetry** or **Traffic analytics** at the fabric level and click **Save**.

     See the editing fabric settings article for your fabric type for more information.

Once **Flow telemetry** or **Traffic analytics** is enabled, the switch will start streaming PTP information to Nexus Dashboard.

4. Review the PTP information provided in the **Telemetry collection status** page.
   a. Click **Manage > Fabrics**.
   b. Locate the fabric that you enabled PTP on and click that fabric.

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

   c. In the **General** area, locate the **Telemetry status** field and click the status entry.

      The **Telemetry collection status** page for that fabric appears, with the **Fabric** tab chosen by default.

      PTP information will be displayed in the fabric's **Telemetry collection status** page, such as the **PTP status** field, which will show **In sync** or **Out of sync**, and the PTP grand leader details, such as the switch in the fabric that is acting as the grand leader.

   d. Click the **Switches** tab, click **... > View details** on a switch, then locate the PTP area to view PTP configuration information at the switch level.

If you configure the fabric with PTP using an internally-elected grand leader clock, you must also enable the ptp clock periodic-update configuration on the grand leader clock. This ensures that the hardware clock stays synchronized with the software clock. If you do not enable this configuration on the grand leader clock, then an anomaly will be triggered, as described below.

Configure PTP: ACI fabrics

Follow these steps to configure PTP in an ACI fabric.

If you plan to use the flow telemetry functions in Nexus Dashboard:

- You must select Telemetry Priority in the ACI fabric node control policy.

In Cisco APIC, choose **Fabric > Fabric Policies > Policies > Monitoring > Fabric Node Controls > <policy-name> > Feature Selection** to select Telemetry Priority. Monitoring <policy-name> should be attached to **Fabric > Fabric Policies > Switches > Leaf/Spine Switches > Profiles**.

- You must also enable Precision Time Protocol (PTP) on Cisco APIC so that Nexus Dashboard can correlate flows from multiple switches accordingly.

  In Cisco APIC, choose **System > System Settings > PTP and Latency Measurement > Admin State** to enable PTP.

  The quality of the time synchronization via PTP depends on the accuracy of the PTP Grandmaster (GM) clock which is the source of the clock, and the accuracy and the number of PTP devices such as ACI switches and IPN devices in between.

  Although a PTP GM device is generally equipped with a GNSS/GPS source to achieve the nanosecond accuracy which is the standard requirement of PTP, microsecond accuracy is sufficient for Nexus Dashboard Insights and its flow telemetry, hence a GNSS/GPS source is typically not required.

  For a single-pod ACI fabric, you can connect your PTP GM via leaf switches. Otherwise, one of the spine switches will be elected as a GM. For a multi-pod ACI fabric, you can connect your PTP GM via leaf switches or via IPN devices. Your IPN devices should be PTP boundary clocks or PTP transparent clocks so that ACI switch nodes can synchronize their clock across pods. To maintain the same degree of accuracy across pods, it is recommended to connect your PTP GM via IPN devices.

  See the section "Precision Time Protocol" in the *Cisco APIC System Management Configuration Guide* for details about PTP connectivity options.

PTP-related anomalies

When either Traffic Analytics or Flow Telemetry are enabled for NX-OS fabrics, Nexus Dashboard monitors for PTP-related anomalies and will take the appropriate action, as described below.

- When significant corrections occur in the fabric across switches, Nexus Dashboard will identify those corrections and will trigger an anomaly alert with the message Ptp Clock High Correction. This alert will automatically clear if no high correction activity is detected during a continuous 20-minute period. The high correction threshold is set to 10000 nanoseconds.
- Nexus Dashboard will also trigger an anomaly if the switch clock and the Nexus Dashboard clock drift out of sync by 60 seconds or more, with the message ND and Switch Clocks Not In Sync. This anomaly will automatically clear once the switch and the Nexus Dashboard clocks remain synchronized for a continuous 20-minute period.
- You must enable the ptp clock periodic-update configuration on the grand leader clock if you configure the fabric with PTP using an internally-elected grand leader clock, as described above. If you do not enable this configuration on the grand leader clock, then an anomaly will be triggered with the message Ptp Clock Missing Config.

For more information, see Detecting Anomalies in Your Nexus Dashboard.

Copyright

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA

https://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883