



Reviewing System Status for Your Nexus Dashboard, Release 4.3.1

Table of Contents

New and changed information	1
Overview	2
Nodes	3
Storage utilization	3
Add nodes	3
Manage secondary nodes	4
Add secondary nodes	4
Delete a secondary node	5
Manage standby nodes	6
Add standby nodes	6
Replace single primary node with standby node	8
Replace two primary nodes with standby node	9
Delete standby nodes	9
Anomalies	10
Scale compliance anomalies	11
View scale compliance anomalies	12
Advisories	13
System advisory notification	15
Metadata support	16
Telemetry	17
Understanding system status	18
Configuring telemetry on Cisco Catalyst devices	19
Telemetry feature support matrix	19
Resources	20
Features	21
Copyright	22

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	Scale compliance monitoring in Nexus Dashboard	Beginning with Nexus Dashboard release 4.3.1, Nexus Dashboard includes Scale Compliance, a health component that evaluates fabric telemetry resource counts against the limits specified in the Cisco Nexus Dashboard Verified Scalability Guide . When a monitored resource exceeds its verified limit, Nexus Dashboard raises a system-level anomaly for the affected cluster. For more information, see Scale compliance anomalies .

Overview

1. Navigate to **Overview** in **System Status**.

Admin > System Status > Overview.

2. Review the information in **Overview**.

Field	Description
Anomaly level	Provides Nexus Dashboard-level anomaly information. Click the Anomaly level tile to navigate directly to the Anomalies tab in System Status . See Anomalies for more information.
Connectivity to Intersight	Provides the status for connectivity to Intersight. Click Setup Intersight to navigate to the Intersight Device Connector area. See Working With Intersight for more information.
Fabrics	Show this information: ▪ The number of fabrics currently onboarded in your Nexus Dashboard ▪ The connectivity status of those fabrics ▪ The fabric types of all the fabrics in your Nexus Dashboard ▪ The license tiers used by the fabrics in your Nexus Dashboard Click View all to navigate directly to Manage > Fabrics.
Cluster nodes	Provides information on the nodes that are currently part of the cluster and the health status for those nodes. Click View all to navigate directly to Nodes in System Status . See Nodes for more information.

Nodes

1. Navigate to **Nodes** in **System Status**.

Admin > System Status > Nodes.

2. Review the information provided in **Nodes**.

Field	Description
Utilization	Provides utilization information for CPU, memory, and storage. See Storage utilization for more information.
Nodes by status	Provides additional information on each node in the cluster. Click Add node to add a node to the cluster. See Add nodes for more information.

Storage utilization

The **Storage** utilization value is the average allocated storage percentage of all nodes in the Nexus Dashboard cluster. It reflects disk-level allocation, not filesystem usage. This value is expected to remain stable and does not increase or decrease unless a feature or service in the cluster changes.

A high storage utilization value does not necessarily indicate a cluster issue. To verify cluster health, use the **acs health** command.

Add nodes

1. In **Nodes** in **System Status**, click **Actions > Add node**.

The **Add node** page opens.

2. In the **Deployment details** area, provide the credentials information for the node, then click **Validate**.
 - o For physical nodes, this is the IP address, username, and password of the server's CIMC. The CIMC will be used to configure the rest of the information on the node.
 - o For virtual nodes, this is the IP address and rescue-user password you defined for the node when deploying it.
3. In the **General** area:
 - a. Provide the name and serial number of the node.
 - b. From the **Type** dropdown, choose **Secondary**.
4. In the **Management network** area, provide the management network information.
 - o For physical nodes, you must provide the management network IP address, netmask, and gateway now.
 - o For virtual nodes, the management network information will be pre-populated with the information pulled from the node based on the IP address and credentials you provided in the previous sub-step.
5. In the **Data network** area, provide the data network information.

You must provide the data network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

6. (Optional) Provide IPv6 information for the management and data networks.
 - Nexus Dashboard supports dual stack IPv4/IPv6 for the management and data networks.
 - If you want to provide IPv6 information, you must do it when adding the node.
 - All nodes in the cluster must be configured with either only IPv4 or dual IPv4/IPv6 stack.
7. In the **Enable BGP** field, click the toggle to enable this feature, if necessary.
8. Click Save to add the node.

The configuration will be pushed to the node and the node will be added to the list in the GUI.

Manage secondary nodes

You can add a number of secondary nodes to an existing 3-node cluster for horizontal scaling to enable application co-hosting.



- Secondary nodes are not supported for cloud form factors of Nexus Dashboard clusters deployed in AWS or Azure.
- Secondary nodes are qualified for IPFM fabric types. For more information about IPFM fabrics, see [Editing IP Fabric for Media \(IPFM\) Fabric Settings](#).

Add secondary nodes

This section describes how to add a secondary node to your cluster to enable horizontal scaling.

Before you begin

- Ensure that the existing primary nodes and the cluster are healthy.
- Prepare and deploy the new node.
- Ensure that the node you are adding is powered on.
- If you are adding a physical node, ensure that you have the new node's CIMC IP address and login information.

You will need to use the CIMC information to add the new node using the Nexus Dashboard GUI.

- If you are adding a virtual node, ensure that you have the node's management IP address and login information.

To add a secondary node:

1. Navigate to **Nodes** in **System Status**.
Admin > System Status > Nodes.
2. In the main pane, click **Actions > Add node**.

The **Add Node** page opens.

3. In the **Add Node** screen, provide the node information.

- a. In the **Deployment details** area, provide the credentials information for the node, then click **Validate**.
 - For physical nodes, this is the IP address, username, and password of the server's CIMC. The CIMC will be used to configure the rest of the information on the node.
 - For virtual nodes, this is the IP address and **rescue-user** password you defined for the node when deploying it.
- b. In the **General** area, provide the name and serial number of the node.
- c. From the **Type** dropdown, choose **Secondary**.
- d. Provide the **Management network** information.

For virtual nodes, the management network information will be pre-populated with the information pulled from the node based on the IP address and credentials you provided in the previous sub-step.

For physical nodes, you must provide the management network IP address, netmask, and gateway now.

- e. Provide the **Data network** information.

You must provide the data network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

- f. (Optional) Provide IPv6 information for the management and data networks.

Nexus Dashboard supports dual stack IPv4/IPv6 for the management and data networks.

If you want to provide IPv6 information, you must do it when adding the node.

All nodes in the cluster must be configured with either only IPv4 or dual IPv4/IPv6 stack.

- g. In the **Enable BGP** field, click the toggle to enable this feature, if necessary.

4. Click **Save** to add the node.

The configuration will be pushed to the node and the node will be added to the list in the GUI.

Delete a secondary node

Before deleting a secondary node:

- Ensure that the primary nodes and the cluster are healthy
- Check for any anomalies
- Delete only one secondary node at a time. After deleting a secondary node, wait until the cluster returns to a healthy state before deleting another secondary node. You can verify the cluster health from **Admin > System Status > Overview** or by running the `acs health` command. Proceed with the next secondary node deletion only after the cluster health is healthy.

See [Overview](#) for more information.

To delete an existing secondary node:

1. Navigate to **Nodes** in **System Status**.

Admin > System Status > Nodes.

2. Select the checkbox next to the secondary node you want to delete.
3. From the **Actions** menu, choose **Delete** to delete the node.
4. Verify that the cluster is healthy before you delete another secondary node.

Manage standby nodes

You can add up to two standby nodes, which you can use to quickly restore the cluster functionality in case one or more primary nodes fail by replacing the failed primary node with the standby node.

Standby nodes are similar to secondary nodes in deployment, initial configuration, and upgrades. However, unlike secondary nodes, the cluster will not use the standby nodes for any workloads.



Standby nodes are not supported for single-node clusters or clusters deployed in AWS or Azure.

The following two cases are supported:

- Single primary node failure

You can use the UI to convert the standby node into a new primary node.

- Two primary nodes failure

You will need to perform manual failover of one of the nodes to restore cluster functionality. Then fail over the second node using standard procedure.

Add standby nodes

This section describes how to add a standby node to your cluster for easy cluster recover in case of a primary node failure.

Before you begin

- Ensure that the existing primary nodes and the cluster are healthy.
- Prepare and deploy the new node.

You can failover only between nodes of identical types (physical or virtual), so you must deploy the same type of node as the nodes in your cluster which you may need to replace. In case of virtual nodes deployed in VMware ESX, which have two node profiles (**OVA-app** and **OVA-data**), you can failover only between nodes of the same profile.

- Ensure that the node you are adding is powered on.
- If you are adding a physical node, ensure that you have the new node's CIMC IP address and login information.

You will need to use the CIMC information to add the new node using the Nexus Dashboard GUI.

- If you are adding a virtual node, ensure that you have the node's management IP address and login information.

To add a standby node:

1. Navigate to **Nodes** in **System Status**.

Admin > System Status > Nodes.

2. In the main pane, click **Actions > Add node**.

The **Add Node** page opens.

3. In the **Add Node** screen, provide the node information.

- a. In the **Deployment details** area, provide the credentials information for the node, then click **Validate**.

- For physical nodes, this is the IP address, username, and password of the server's CIMC. The CIMC will be used to configure the rest of the information on the node.
- For virtual nodes, this is the IP address and **rescue-user** password you defined for the node when deploying it.

- b. In the **General** area, provide the name and serial number of the node.

- c. From the **Type** dropdown, select **Standby**.

- d. Provide the **Management network** information.

For virtual nodes, the management network information will be pre-populated with the information pulled from the node based on the IP address and credentials you provided in the previous sub-step.

For physical nodes, you must provide the management network IP address, netmask, and gateway now.

- e. Provide the **Data network** information.

You must provide the data network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

- f. (Optional) Provide IPv6 information for the management and data networks.

Nexus Dashboard supports dual stack IPv4/IPv6 for the management and data networks.

If you want to provide IPv6 information, you must do it when adding the node.

All nodes in the cluster must be configured with either only IPv4 or dual IPv4/IPv6 stack.

- g. In the **Enable BGP** field, click the toggle to enable this feature, if necessary.

4. Click **Save** to add the node.

The configuration will be pushed to the node and the node will be added to the list in the GUI.

Replace single primary node with standby node

This section describes failover using a pre-configured **standby** node. If your cluster does not have a standby node, follow the steps described in one of the sections in [Cisco Nexus Dashboard Troubleshooting](#) instead.

Before you begin

- Ensure that at least 2 primary nodes are healthy.
- Ensure that you have at least one **standby** node available in the cluster.

Setting up and configuring **standby** nodes is described in [Add standby nodes](#).

- Ensure that the **primary** node you want to replace is powered off.



You cannot re-add the **primary** node you are replacing back to the cluster after the failover is complete. If the **primary** node you replace is still functional and you want to re-add it to the cluster after the failover, you must factory reset or re-image it as described in [Cisco Nexus Dashboard Troubleshooting](#) and add it as a **standby** or **primary** node only.

To failover a single primary node:

1. Navigate to **Nodes** in **System Status**.

Admin > System Status > Nodes.

2. Click the **Actions** (...) menu next to the **Inactive** primary node that you want to replace.
3. Choose **Failover**.

Note that you must have a standby node already configured and added or the **Failover** menu option will not be available.

4. In the **Fail Over** window that opens, select a standby node from the dropdown.



You can promote a **standby** node to a **primary** node only when the **standby** node's status is **Inactive**.

5. Click **Save** to complete the failover.

The failed primary node will be removed from the list and replaced by the standby node you selected. The status will remain **Inactive** while the services are being restored to the new primary node.

It can take up to 10 minutes for all services to be restored, at which point the new primary node's status will change to **Active**.

If you want to add the old node back to the Nexus Dashboard cluster as a **standby** node, you must first reset it to factory defaults by using the **acs reboot factory-reset** command and then add it back by using the procedure described in [Add standby nodes](#).

Replace two primary nodes with standby node

The option to replace two primary nodes with a standby node is not supported. Instead, if one cluster becomes unavailable, you will recover that cluster from a backup that is available on another cluster. See the section "Perform a dynamic recovery on a cluster" in [Nexus Dashboard Troubleshooting](#) for more information.

Delete standby nodes

Before deleting a standby node:

- Ensure that the primary nodes and the cluster are healthy
- Check for any anomalies

See [Overview](#) for more information.

To delete an existing standby node:

1. Navigate to **Nodes** in **System Status**.

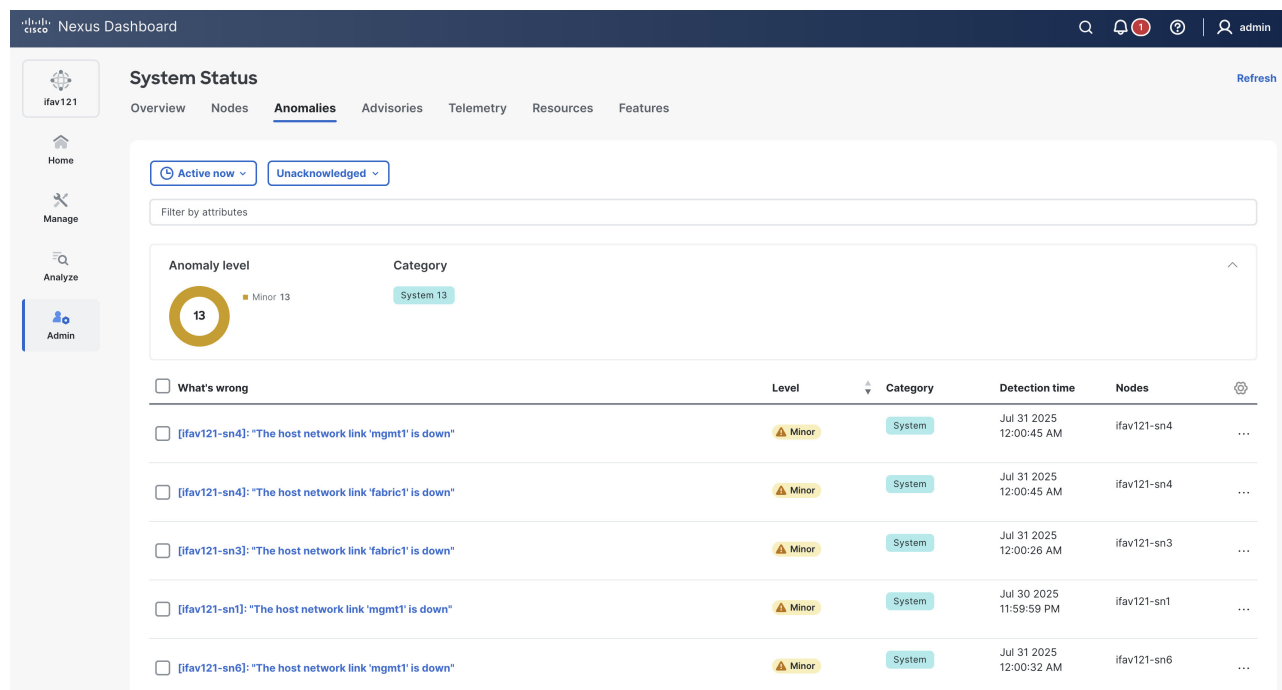
Admin > System Status > Nodes.

2. Select the checkbox next to the standby node you want to delete.
3. From the **Actions** menu, choose **Delete** to delete the node.

Anomalies

The **Anomalies** tab allows you to quickly monitor platform-level anomalies detected on Nexus Dashboard. It highlights critical, high-severity events that require your prompt attention and resolution to keep the system healthy and stable.

1. To navigate to the **Anomalies** tab in the **System Status** page, choose **Admin > System Status > Anomalies**.
2. Review the information provided in the **Anomalies** table.



The screenshot displays the 'System Status' page with the 'Anomalies' tab selected. The left sidebar shows navigation options: Home, Manage, Analyze, and Admin. The main content area has a 'System Status' header with tabs for Overview, Nodes, Anomalies (active), Advisories, Telemetry, Resources, and Features. Below the header, there are filters for 'Active now' and 'Unacknowledged'. A summary card shows 'Anomaly level' as 13 Minor and 'Category' as System 13. Below this is a table of anomalies.

What's wrong	Level	Category	Detection time	Nodes
☐ [ifav121-sn4]: "The host network link 'mgmt1' is down"	Minor	System	Jul 31 2025 12:00:45 AM	ifav121-sn4
☐ [ifav121-sn4]: "The host network link 'fabric1' is down"	Minor	System	Jul 31 2025 12:00:45 AM	ifav121-sn4
☐ [ifav121-sn3]: "The host network link 'fabric1' is down"	Minor	System	Jul 31 2025 12:00:26 AM	ifav121-sn3
☐ [ifav121-sn1]: "The host network link 'mgmt1' is down"	Minor	System	Jul 30 2025 11:59:59 PM	ifav121-sn1
☐ [ifav121-sn6]: "The host network link 'mgmt1' is down"	Minor	System	Jul 31 2025 12:00:32 AM	ifav121-sn6

The **Anomalies** table displays filtered anomalies. By default, the anomalies are sorted by level. Click the column heading to sort the anomalies in the table. To view cleared anomalies, apply a filter to list anomalies from the past 15 minutes (or a similar recent time frame), and configure the table to show the **Status** column, where the status is displayed as either **Active** or **Cleared**. An **Active** status means the anomaly is present in your network, while a **Cleared** status means the anomaly is no longer present.

3. Click an anomaly to view more information.

The **Anomaly name** page displays the following details.

- **What's wrong?** – provides a problem description with the specific affected objects.
- **What triggered this anomaly?** – provides the primary source of the anomaly.
- **What's the impact?** – explains the potential impact if the problem is not fixed.
- **How do I fix it?** – provides prescriptive recommendations.

Nodes Network Device Down

What's wrong?
[ifav121-sn4]: "The host network link 'mgmt1' is down"

Anomaly level ▲ Minor

Status active ▲ Last seen
Last seen: Jul 31, 2025, 08:34:06 AM

Category	Nodes	Initial detection time	Tags	Verification status	Assigned to
System	ifav121-sn4	Jul 31 2025 12:00:45 AM	-	New	-

What triggered this anomaly?
Node level networkLinkStatus anomaly is triggered if one or more host level network links goes down.

What's the impact?

- Network stability is fundamental for services to remain operational. Network links going down or flapping impacts several aspects including:
 - Service disruptions: Pods may lose connectivity to other pods or external services.
 - Failed health checks: Kubernetes may evict healthy pods if they fail readiness/liveness checks due to network issues.
 - Cluster instability: Control plane nodes may struggle to reach kubelets or etcd.
 - Data corruption: if NAS is used, dropped links can cause write failures or timeouts.
 - Management link failures might block all the management operations via SSH or UI.

Persistent link flaps or link failures may lead to total outage.

How do I fix it?

Recommended solution

- You should always connect redundant links for both data and management

For more information, see [Detecting Anomalies in Your Nexus Dashboard](#).

Scale compliance anomalies

Beginning with Nexus Dashboard release 4.3.1, scale compliance anomalies appear in Nexus Dashboard. Scale compliance checks monitored resource counts, including fabric telemetry, Controller, and Orchestrator resources, against the verified scale limits for the cluster deployment profile and form factor. When a monitored resource reaches its verified limit, Nexus Dashboard raises one Major system anomaly for the affected cluster.

Nexus Dashboard periodically checks scale compliance and updates the anomaly while any monitored resource remains above its verified limit. The anomaly is cleared after the affected values fall below their verified scale limits.

Depending on the enabled services and fabric types, scale compliance can monitor the following resources:

- Switches and endpoints
- Physical interfaces
- IPv4 routes, IPv6 routes, VRFs, and fabrics
- Layer 2 VNIs, Layer 3 VNIs, multicast routes, MAC entries, VLANs, and VTEPs
- Flow and Traffic Analytics scale conditions

Scale limits can vary by Nexus Dashboard deployment profile and form factor, including both physical and virtual deployments. Nexus Dashboard evaluates each deployment against the verified limits for its detected profile. To check the limits at which a scale compliance anomaly is raised for a particular profile, refer to the [Cisco Nexus Dashboard Verified Scalability Guides](#) for your release.

When a cluster contains both ACI and NDFC fabrics, Nexus Dashboard uses the lower applicable scale limit for each monitored resource. In federated deployments, each cluster is evaluated independently, and scale compliance anomalies appear only on the **System Status** page of the affected cluster.

When Nexus Dashboard generates a scale compliance anomaly, the **Notifications** bell icon indicates that a new system anomaly is available. Open the notification to go to the system anomaly details, where you can view the affected parameter, current usage, verified limit, and percentage of use. For example, the anomaly details might show **Switches: current usage 513 of limit 500 (103%)**.

If Nexus Dashboard cannot evaluate a scale parameter because of an API failure, it reports a system error fault for that parameter.

After a restore, Nexus Dashboard evaluates scale compliance during the next evaluation cycle and raises a scale compliance anomaly only if the current values meet the anomaly conditions. No operator action is required.

View scale compliance anomalies

Follow these steps to view scale compliance anomalies.

1. To navigate to the **System Status** page, choose **Admin > System Status > Anomalies**.
2. Review the information provided in the table.

Scale compliance anomalies are listed with the category **System**. You can also view system anomalies from **Analyze > Anomalies** by enabling **Include system anomalies**.

3. Click a scale compliance anomaly to view more information.

The **Anomaly** page displays the following details:

- **What's wrong?**—Lists the monitored resources that have reached their verified limits. The details include the current usage, verified limit, and percentage of use for each affected resource.
- **What triggered this anomaly?**—Describes which verified limit was reached and the deployment profile used for evaluation.
- **What's the impact?**—Explains potential effects such as telemetry delays, reduced visibility, degraded query performance, or service instability.
- **How do I fix it?**—Provides remediation guidance, such as reducing the monitored object count, disabling telemetry for selected fabrics, distributing fabrics across clusters, or moving to a larger supported deployment profile.

Advisories

Nexus Dashboard uses metadata bundles to detect field notices, software and hardware end-of-life (EoL) and end-of-sale (EoS) announcements, as well as PSIRTs that affect the network cluster nodes, and it generates advisories. These advisories recommend actions to support your network for optimal performance. Previously, Nexus Dashboard limited advisories for fabric nodes. With this release, Nexus Dashboard includes advisories for cluster nodes. For more information, see [Metadata support](#) and [Identifying Advisories in Your Nexus Dashboard](#).

Follow these steps to view the cluster advisories.

1. Navigate to **Advisories** tab in the **System Status** page.

Go to **Admin > System Status > Advisories**.

2. Review the information provided in the **Advisories** table.

System Status Refresh

Overview Nodes Anomalies **Advisories** Telemetry Resources Features

Active now Acknowledged

Filter by attributes

Level

- Critical: 2
- Major: 2
- Warning: 4

Category

- Hardware EOL: 3
- Software EOL: 1
- PSIRT: 3
- Field notice: 1

Title	Advisory level	Category	Nodes
End-of-Sale and End-of-Life Announcement for the Cisco SE-NODE-G2	Warning	Hardware EOL	ifav121-sn1 ifav121-sn2 ifav121-sn3 ZB
End-of-Sale and End-of-Life Announcement for the Cisco N9K-C9236C and N9K-C92304QC	Critical	Hardware EOL	ifav121-sn4 ifav121-sn5 ifav121-sn6
End-of-Sale and End-of-Life Announcement for the Cisco APIC-M3, APIC-L3 and SE-NODE-G2 - Cisco	Warning	Hardware EOL	ifav121-sn3 ifav121-sn1 ifav121-sn2 ZB
CSCwi62525: Vulnerabilities in paramiko 2.6.0 CVE-2023-48795 and others	Warning	PSIRT	ifav121-sn4 ifav121-sn3 ifav121-sn1 ifav121-sn2 ifav121-sn5 ZB ifav121-sn6



Use the **Filter** drop-down list to choose the appropriate column name filter and display specific columns in the **Advisories** table.

The **Advisories** table displays the following information.

Field	Description
Title	Displays the name of the advisory as published by Cisco PSIRT or other sources.
Advisory level	Specifies if the custom dashboard is shared or private.
Category	Specifies the type of advisory, such as PSIRT, Field notice, Hardware, and Software EoL.
Nodes	Displays advisories for specific nodes.
What's wrong	Displays advisories of a specific affected object.

Field	Description
Detection time	Display advisories with a specific detection time.
Last seen time	Displays only advisories with a specific last seen time. The last seen time indicates when the advisory was updated while it was active. If the Nexus Dashboard does not clear the advisory status, it stays active.



Nexus Dashboard lists nodes on the **Advisories** page for 30 minutes after removing them from the cluster. This intentional delay ensures proper handling of any ongoing processes or alerts related to the nodes before removing them from the **Advisories** table. Hence, the main **Advisories** page displays the nodes that are already removed from the cluster but are still within the 30-minute grace period. However, when you click on a specific advisory to view detailed information, Nexus Dashboard displays a more accurate list of nodes, excluding those that have been removed.

3. Navigate to **Active now > Time Selection**, to choose the date and time range.

By default, **Active now** is chosen. You can customize the date and time range to determine the advisories data displayed in the **Advisories** table.

4. Choose the advisories from the **Advisories** table and click **Acknowledge advisories** to acknowledge advisories.

The screenshot displays the Nexus Dashboard's Advisories page. The interface includes a sidebar with navigation links (Home, Manage, Analyze, Admin) and a main content area. The 'System Status' section is active, showing various tabs. The 'Advisories' tab is selected, displaying a summary of advisories by level (Critical, Major, Warning) and category (Hardware EOL, Software EOL, PSIRT, Field notice). A table below lists specific advisories, including their titles, levels, categories, and associated nodes. One advisory is highlighted with a blue background and a checkmark in the selection column.

5. By default, all the unacknowledged advisories are displayed in the **Advisories** table. Once you acknowledge an advisory, choose **Acknowledged** from the drop-down list to view all the acknowledged advisories.
6. The **Advisories** page displays the total number of advisories by severity such as **Critical**, **Major**, and **Warning** in the donut chart under **Level**. You can view the type of advisory, such as PSIRT, field notices, hardware, and software EoL under **Category**.
7. Click an advisory to view more information.

The **Advisory name** page displays these details.

- **What's wrong?** – provides a problem description with the specific affected objects.
- **How do I fix it?** – provides prescriptive recommendations.

The screenshot shows the Cisco Nexus Dashboard interface. The top navigation bar includes the Cisco logo, 'Nexus Dashboard', search, notifications, and user profile. The left sidebar has icons for Home, Manage, Analyze, and Admin. The main content area displays an advisory titled 'End-of-Sale and End-of-Life Announcement for the Cisco N9K-C9236C and N9K-C92304QC'. It includes a 'What's wrong?' section with a link to the Cisco announcement, an 'Advisory level' of 'Critical', and a 'Status active' indicator. Below this is a table titled 'What's the impact?' with columns for Name, Model, and Serial number. The table lists three items: #av121-sn4, #av121-sn5, and #av121-sn6, all with model ND-NODE-L4T. At the bottom, there is a 'How do I fix it?' section with a 'Recommended solution' link.

Name	Model	Serial number
#av121-sn4	ND-NODE-L4T	WZP28170ZSS
#av121-sn5	ND-NODE-L4T	WZP28170ZSK
#av121-sn6	ND-NODE-L4T	WZP28179JRH



You can also view the cluster advisories by navigating to **Analyze > Advisories**. Click the **Include system advisories** toggle button to include system advisories. For more information, see [Identifying Advisories in Your Nexus Dashboard](#).

System advisory notification

When there is an active system advisory, a notification alert appears on the **Notifications** bell icon located in the common navigation bar at the top of the page. Click the notification bell icon to open the **Notifications** pane. In the **Notifications** pane, click **View system advisories**. Nexus Dashboard redirects you to the **System Status** page, where you can review the full list of current and past system advisories in the **Advisories** table.

System Status

Overview Nodes Anomalies **Advisories** Telemetry Resources

Active now Unacknowledged

Filter by attributes

Level: Warning 1 Category: Hardware EOL 1

Title	Advisory level	Ca
End-of-Sale and End-of-Life Announcement for the Cisco APIC-M3, APIC-L3 and SE-NODE-G2 - Cisco	Warning	H

Notifications

System advisories

There are advisories that are applicable to one or more nodes within this cluster. Please click the link for more details.

[View system advisories](#)

Metadata support

Nexus Dashboard uses metadata bundles to detect latest bug signatures, PSIRTs, field notices, and end-of-life notices. Cisco Intersight Cloud regularly updates, validates, and makes metadata packages available. Nexus Dashboard connects to the [Cisco Intersight Cloud](#) via an embedded device connector, which periodically retrieves the updated metadata packages. For air-gapped environments, where the Nexus Dashboard is not connected to the Cisco Intersight Cloud, you can securely and manually upload the latest metadata. You can download the bundle updates from [Cisco Intersight](#).

Follow these steps to check the metadata version.

1. Navigate to **Admin > System Settings**.
2. In the **General** tab, under **Metadata** you can view the metadata version.
3. Click **Edit** to update the metadata version.

The **Metadata** page appears. In the **Update metadata version** area, you can manually upload the latest metadata files.

Telemetry

Beginning with Nexus Dashboard release 4.3.1, the **Telemetry** status page provides accurate reporting of Traffic analytics capabilities for switches within ACI fabrics.

The system now evaluates specific hardware models and software versions to determine the supported traffic analytics mode. This enhancement ensures that the displayed status reflects the actual capabilities of the hardware, correcting previous issues where switches were incorrectly labeled as "Traffic analytics full mode capable" when they only supported "Compatibility mode."

The **TA capability** column is updated with improved logic to provide precise information regarding feature mode support for each switch in your fabric.

You can monitor the number of switches that support Full, Compatibility, and Unsupported modes from the following pages:

- Navigate to **System Status Telemetry > Fabrics**.
- Navigate to **System Status Telemetry > Switches**.

Beginning with Nexus Dashboard 4.2.1, Nexus Dashboard collects basic telemetry data from Cisco Catalyst 9000 (CAT9K) devices running IOS-XE 17.15. Data collected from legacy protocols such as SNMP, Syslog, Netflow, and sflow provide basic telemetry and analytics information to represent device health and monitoring.


1. Navigate to **Telemetry** in **System Status**.

Admin > System Status > Telemetry.

- Click the **Fabrics** tab to view telemetry status information for the fabrics in your Nexus Dashboard.
- In Nexus Dashboard 4.3.1, the **TA capability** column is enhanced to display the number of switches in a fabric supporting Traffic analytics full, Traffic analytics compatibility, and Unsupported.

For example, in the following screenshot, the TA capability column shows the number of switches in each mode. For example, a value of 3/1/2 means:

- 3 switches are Traffic analytics full mode capable
- 1 switch is Traffic analytics compatibility mode capable
- 2 switches are unsupported (do not support Traffic analytics capability because of hardware limitation or software version)



The screenshot shows the 'System Status' page with the 'Telemetry' tab selected. Under 'Telemetry status', the 'Fabrics' sub-tab is active. A table lists the fabric 'DC-blr9'. The 'TA capability' column for this fabric displays '3 / 1 / 2', indicating 3 switches in full mode, 1 in compatibility mode, and 2 unsupported.

Fabrics	TA mode	TA capability	Telemetry config status	Software telemetry	Flow collection	Software telemetry switch status	Flow collection switch status
DC-blr9	Disabled	3 / 1 / 2	OK	Enabled	Disabled	-	-

- Click the **Switches** tab to view telemetry status information for the fabrics in your Nexus Dashboard. In Nexus Dashboard 4.3.1, the **TA capability** column displays the traffic analytics capability for each switch as Full, Compatibility, and Unsupported.



After upgrading to Nexus Dashboard 4.1.1, if any fabric shows a **Telemetry** configuration status as **Pending updates**, all other configurations or runtime states (at both fabric and switch level) should be ignored until a redeploy is triggered.

Understanding system status

Nexus Dashboard processes your fabric's telemetry through different jobs, services and tasks that reflect what you see on screen. The statuses are summarized at the fabric level. Following is a brief description for each of them:

- **Assurance:** Indicates the status of the last assurance collection job. Hovering over this field will detail when assurance was last run, and whether it was a scheduled or on-demand task.
- **Capacity:** Ensures all validated switch capacity limitations are in conformance. It is expected for controllers to report **No Data** as this is not applicable to these devices.
- **Hardware resources:** Monitors the health of all switch HW resources including CPU, memory, fans, power supplies, storage and environmental levels are healthy.
- **Statistics:** Indicates the status of the collection of switch and interface level metrics. This collection is refreshed every 5 minutes.
- **Endpoints:** Displays the collection status of a switches Endpoint records. Hovering over this field will display the last update timestamp. Endpoint collection is not applicable to controllers or spines switches that have no endpoints connected.
- **Bug scan:** Provides the status of the previous bug scan analysis. Hovering over this field will provide the timestamp of the last run attempt.
- **Best practices:** Displays the status of the last best practices scan. Hovering over this field will provide the timestamp of the last run attempt.
- **Telemetry collection status:** This is available only at telemetry status in **Fabric Overview (Manage > Fabrics > FabricXYZ)**. This is dynamically updated and represents the data streaming status.
- **Telemetry configuration status:** Indicates that basic telemetry has been enabled on devices in the fabric. Status are:
 - **OK:** All switches have been successfully configured for telemetry streaming.
 - **Not OK:** Telemetry configuration for all switches in the fabric has failed or pending change control change.
 - **Partial OK:** Some switches have been successfully configured for telemetry streaming, some failed.
 - **In progress:** Telemetry configuration attempting to change state (Telemetry Pause/Resume, Pending change control).
 - **Pending updates:** Indicates new telemetry configurations are available, can be availed with 'redeploy' action.
 - **Out of sync :**Indicates a restore operation is in-progress, should be completed with

'reconfigure' action.

- **Software telemetry status:** Displays the software telemetry status for each switch. The value for this property will show **Enabled**, **Disabled** or a **Pending** state (if Change Control is enabled).
- **Flow collection:** Indicates flow collection configuration status.

Configuring telemetry on Cisco Catalyst devices

Follow these steps to configure telemetry on Cisco Catalyst 9000 devices running IOS-XE 17.15:

1. Configure Telemetry Receiver on CAT9K

```
CAT9K-1#show telemetry receiver name nexusdb
```

2. Configure Telemetry Subscription

Based on the filter specified, required data is collected. Subscriptions can be created for MOs such as endpoint, MAC, ARP, and interface.

```
CAT9K-1# show telemetry ietf subscription 303 detail
```

3. Configure UTR-PL for telemetry processing

This enables telemetry data to be received and parsed before being forwarded to downstream services for processing.

Telemetry feature support matrix

Category	Features
Inventory	Switches, controllers, interfaces, and endpoints with historical data
Hardware resources	CPU, memory, power, temperature, and interfaces with historical information
Topology	Fabric-level with switch, interface, and inventory view for endpoint visibility
Traffic analytics	Fabrics can be either in TA full or TA compatible mode. Mixed mode is not supported where some switches within the fabric are either in TA full and TA compatible mode. NOTE: For TA compatibility mode, navigate to Manage > Fabric, select the fabric and click Actions > Edit fabric settings. Go to the Telemetry tab, select Traffic analytics compatibility mode and click Save.

Resources

Resources provides real-time information about the resource utilization of your Nexus Dashboard cluster.

1. Navigate to **Resources** in **System Status**.

Admin > System Status > Resources.



In Nexus Dashboard 4.3.1, the digital news feed button is removed from the **Admin > System Status > Resources** page.

Features

Features provides information on the features that are enabled in your Nexus Dashboard and the health status for those features.

1. Navigate to **Features** in **System Status**.

Admin > System Status > Features.

Copyright

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

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

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

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

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

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

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

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

Americas Headquarters

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

<https://www.cisco.com>

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883