



Cisco Nexus Dashboard Insights
Analysis Hub, Release 6.4.1 - For
Cisco ACI

Table of Contents

New and Changed Information	2
Sustainability Report	3
Sustainability Report	3
View Sustainability Report for Switches	4
View Sustainability Report for PDUs	6
Conformance Report	8
Conformance Report	8
Access Conformance Report	9
View Conformance Report	9
Connectivity Analysis	11
Connectivity Analysis	11
Guidelines and Limitations	11
Schedule a Connectivity Analysis	12
Connectivity Analysis Dashboard	13
Connectivity Analysis Error Scenarios	17
Filtering Information	18
Delta Analysis	19
Delta Analysis	19
Guidelines and Limitations	19
Create Delta Analysis	20
View Delta Analysis	20
View Health Delta Analysis	21
View Policy Delta Analysis	22
Log Collector	25
Log Collector	25
Log Collector Dashboard	25
TAC Initiated Log Collector	26
Upload logs to Cisco Intersight Cloud	26
Compliance	28
Compliance	28
Guidelines and Limitations	29
Create Compliance Communication Rule	30
Create Compliance Rule with Snapshot Selection	33
Create Import Configuration Compliance	33
Template Based Compliance	34
Create Compliance Rule with Manual Configuration	44
Matching Criteria	47
Compliance Rules	49
Compliance Analysis	50
Policy CAM	52
About Policy CAM	52

Pre-Change	57
Pre-Change Analysis	57
Pre-Change Analysis Options	59
Guidelines and Limitations	60
Support for Multiple Objects in Pre-Change Analysis	61
Known Issues for Pre-Change Analysis	62
Create Pre-Change Analysis Job	62
Download Pre-Change Analysis Job	63
Copyright	65

First Published: 2024-02-16

Last Modified: 2024-04-01

Americas Headquarters

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

<http://www.cisco.com>

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883

New and Changed Information

The following table provides an overview of the significant changes up to the current release. The table does not provide an exhaustive list of all changes or the new features up to this release.

New Features and Changed Behavior in the Cisco Nexus Dashboard Insights

Feature	Description	Release	Where Documented
Sustainability Report for PDUs	If you have Panduit PDUs onboarded to Nexus Dashboard, Sustainability Report for PDUs gives you insight into how much electricity your devices and/or Panduit PDUs are using.	6.4.1	Sustainability Report
Compliance rule description enhancement	This feature allows you to add a custom description when a compliance rule is created.	6.4.1	Create Compliance Communication Rule
Compliance analysis user interface enhancement	Compliance analysis UI has enhanced the Anomalies from Violations table to provide more detailed information.	6.4.1	Compliance Analysis
Compliance Rules navigation change	Compliance Rules are now accessible from Manage > Rules > Compliance Rules .	6.4.1	Create Compliance Communication Rule
Delta Analysis report enhancement	The venn diagram in the Delta Analysis report page has been replaced with a more clear and concise user interface.	6.4.1	View Delta Analysis

This document is available from your Nexus Dashboard Insights GUI as well as online at www.cisco.com. For the latest version of this document, visit [Cisco Nexus Dashboard Insights Documentation](#).

Sustainability Report

Sustainability Report

Cisco Nexus Dashboard Insights Sustainability report helps you monitor, predict, and improve your network's energy usage, its related carbon emissions, and its total energy cost. Sustainability report enables you to get insights on energy utilization, CO2 emissions, and energy cost for all your sites on a monthly basis.

The report is generated by calculating the monthly values for Power Consumption and by summing the usage data across all of your devices at each of your sites for every single day in the selected month. This data is then combined with our third-party services such as Electricity Maps to provide greater insight into what that usage means in terms of energy cost, estimated emissions, and estimated switch power consumption.

The summary area of the report contains information such as estimated cost, estimated switch power consumption, sources of emission, and estimated emissions.

- Estimated Cost gives you insight into any expected increase or decrease in your sites' energy bills based on your monthly energy use.
- Estimated Switch Power Consumption gives you insight into how efficiently your switches are using electricity. Estimated PDU Power Consumption gives you insight into how much electricity your devices and/or Panduit PDUs are using.
- Estimated Emissions gives you insight into the sustainability your sites have on your total CO2 emissions, based on the sources and amount of energy used.

If you have Panduit PDUs onboarded to Nexus Dashboard, you can use the Data Source toggle to see two different electricity values on Sustainability Report - one for Switches only, and one for PDUs.

- Switch Data - Uses only the electricity data reported by individual switches added to a site.
- PDU Data - Uses the electricity data reported by a supported PDU, which could include switches, fans, and any other devices physically plugged into the PDU.

Depending on which value you choose in the Data Source toggle, the values calculated for your other metrics, including estimated cost and emissions, will vary.

Using Sustainability report you can

- Better anticipate increases in your sites' energy bills so that your budgets more accurately reflect real-world usage.
- Better follow the hourly energy usage of an individual site. By spreading out usage to avoid peak hours surcharges, you may be able to lower your electricity bill over time.
- See the direct sustainability impact running your fabric has on climate change. Following your emissions over time also gives you the ability to choose lower-carbon sources and track your progress towards meeting ESG goals.



The retention time for Sustainability Report in Nexus Dashboard Insights is 12

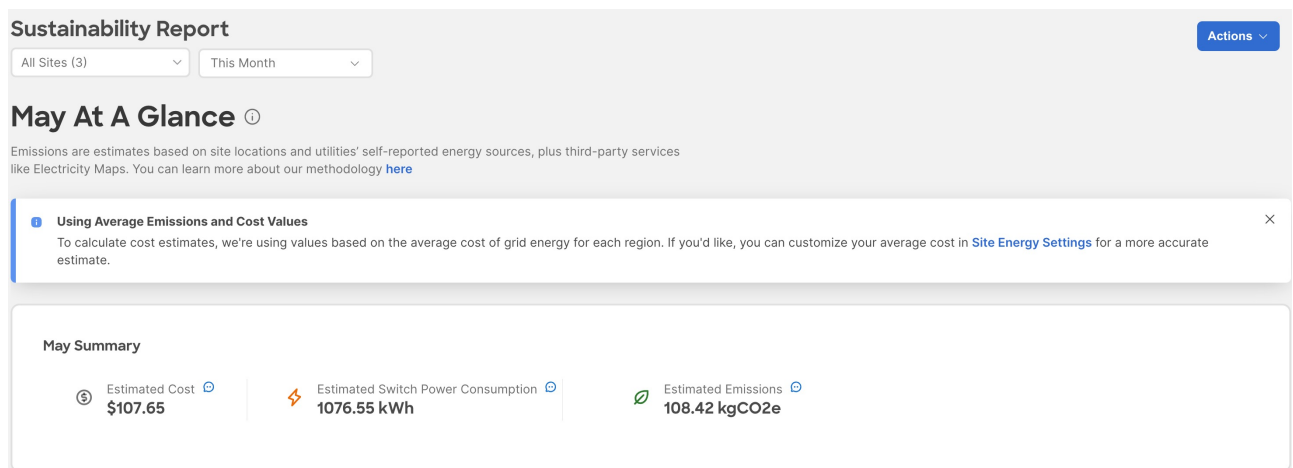
months.

View Sustainability Report for Switches


1. Navigate to **Analyze > Analysis Hub > Sustainability Report**.
2. Select an online site or multiple online sites from the dropdown menu.
3. Select a time range from the dropdown menu.
4. Use the Data Source toggle to display data from switches.
5. Click **Prepare Report**.

Sustainability report displays At A Glance, Cost, Energy, and Emissions information for a particular site in the selected month.

6. The **At A Glance** area displays summary of the estimated cost, estimated switch power consumption, and estimated emissions in the selected month. Click the **Learn More** icon for more information.




7. The **Cost** area displays the estimated daily cost in the selected month and share of daily cost per site.



8. To calculate cost estimates Nexus Dashboard Insights uses values based on the average cost of grid energy for each region. From the Actions menu, select **Site Energy Settings** to customize your average cost for the current month for a more accurate estimate. In the Site Energy Settings page, you can update the cost for a particular site for the current month.
9. The **Energy** area displays the energy usage in the selected month in kWh.

Energy

This month, you've used significantly more energy from the grid across your sites





- The **Emissions** area displays the Total emissions or Efficiency Index per site, estimated monthly carbon dioxide equivalent emissions, average percentage of energy from low-carbon sources and other sources, and percent of total energy used during each three-hour reporting period by source, over all of the days in the selected month.

For Total emissions or Efficiency Index per site use toggle to view the information in graphical format or tabular format.


Emissions

About 71% of your energy this month came from low-carbon sources on average with solar making up the majority



- Select a site from the site dropdown menu to view the Hourly energy mix.

Hourly energy mix displays the amount of total energy used during each three-hour reporting period by source, over all of the days in the selected month. The minimum period before you can generate the next report is 3 hour.




View Sustainability Report for PDUs


1. Navigate to **Analyze > Analysis Hub > Sustainability Report**.
2. Select an online site or multiple online sites from the dropdown menu.
3. Select a time range from the dropdown menu.
4. Use the Data Source toggle to display data from PDUs.
5. Click **Prepare Report**.

Sustainability report displays At A Glance, Cost, Energy, and Emissions information for a particular site in the selected month.

6. The **At A Glance** area displays summary of the estimated cost, estimated switch power consumption, and estimated emissions in the selected month. Click the **Learn More** icon for more information.



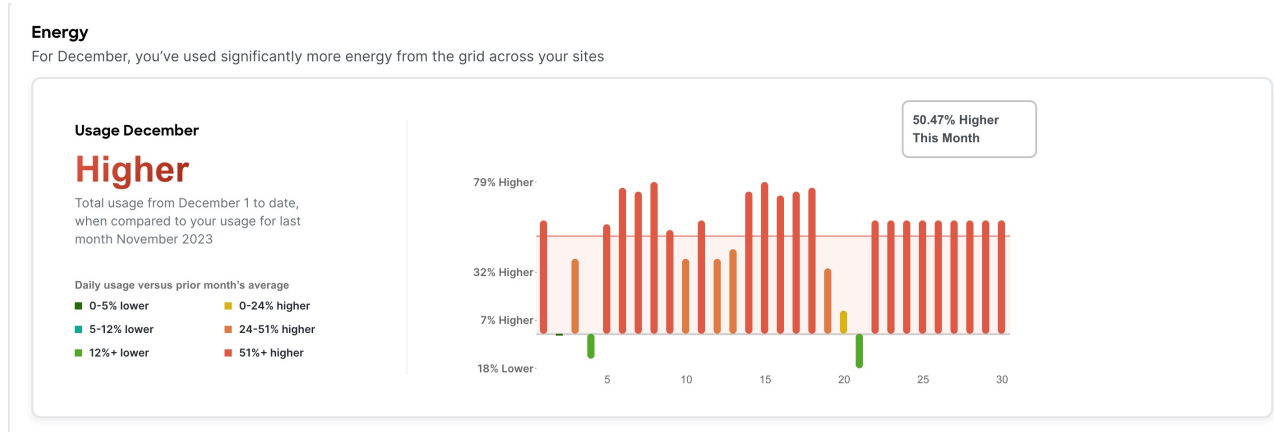
7. The **Cost** area displays the estimated daily cost in the selected month and share of daily cost per site.



8. To calculate cost estimates Nexus Dashboard Insights uses values based on the average cost of grid energy for each region. From the Actions menu, select **Site Energy Settings** to customize your average cost for the current month for a more accurate estimate. In the Site Energy Settings

page, you can update the cost for a particular site for the current month.

9. The **Energy** area displays the energy usage in the selected month in kWh.



10. The **Emissions** area displays the Total emissions per site, estimated monthly carbon dioxide equivalent emissions, average percentage of energy from low-carbon sources and other sources, and percent of total energy used during each three-hour reporting period by source, over all of the days in the selected month.

For Total emissions per site use toggle to view the information in graphical format or tabular format.


Emissions

About 41% of your energy for December came from low-carbon sources on average with nuclear making up the majority



11. Select a site from the site dropdown menu to view the Hourly energy mix.

Hourly energy mix displays the amount of total energy used during each three-hour reporting period by source, over all of the days in the selected month. The minimum period before you can generate the next report is 3 hour.



Conformance Report

Conformance Report

Conformance report enables you to visualize and understand the lifecycle of your hardware and software in the network. This assists you in planning upgrades and hardware refresh. Conformance Report is generated everyday for each site for hardware and software conformance and weekly for each site for scale conformance. In the report you can view the conformance status of software, hardware, combination of both software and hardware, and scale conformance status for sites.

You can use Conformance Report to view current and project the future status of software and hardware inventory in your network against known EoS and EoL notices to ensure conformance. You can also monitor scale conformance status for onboarded sites.

Using Conformance Report you can,

- Minimize risk of running End-of-Sale (EoS) or End-of-Life (EoL) switches.
- View current status of software and hardware inventory in your network against known EoS and EoL notices to ensure conformance.
- Project the future outlook of software and hardware inventory in your network.
- Monitor scale conformance status for onboarded sites.

Conformance Report displays the summary of conformance status for software, hardware, and scale for selected sites.

In the Conformance report, for hardware and software conformance switches are classified into 3 severities based on the software release or hardware platform EoL dates and end of PSIRT dates. The severities include:

- Critical: End of PSIRT date or Last Date of Support occurs in the past.
- Warning: EoL date for software release or EoS for hardware release occurs in the past.
- Healthy: End of PSIRT date, or Last Date of Support and EoL date or software release or EoS for hardware release occurs in the future, or EoL for software release or EoS for hardware release is not announced.

The End of SW Maintenance Releases Date in the End-of-Sale and End-of-Life Announcement and the end of PSIRT date is used as reference milestone to classify the inventory into a category of Critical, Warning, or Healthy.

In the Conformance report, the scale conformance status for sites is based on Cisco's Verified Scalability Guidelines for the software version running in switches and controllers when applicable. The severities include:

- Conformant: All metric values are under 90%.
- Approaching limits: One or more metric values are between 90% and 100%.

- Violated Limits: One or more metric values are over 100%.

Access Conformance Report

Navigate to **Analyze > Analysis Hub > Conformance**.

Select a site from the dropdown menu.

OR

Navigate to **Manage > Sites**.

Select a Site.

In the General section, click **Conformance**.

Click **View Report**.

View Conformance Report

1. Navigate to a Conformance Report. See [Access Conformance Report](#).
2. Choose a site or **All Sites** from the drop-down menu.
3. Choose a current month or a previous month from the drop-down menu. You can choose a previous month only if previous month reports are available.

Conformance Report displays the conformance summary, hardware and software conformance, and scale conformance.

4. The Summary page displays devices by hardware conformance status, devices by software conformance status and scale conformance status for sites or switches. Click **View Conformance Criteria** to learn more.
5. The Hardware or Software page displays conformance status, conformance outlook, and device details.
 - a. In the Conformance Outlook section, click **Overall**, or **Software**, or **Hardware** to view the conformance for software and hardware, software only or hardware only.
 - b. The Device Details lists details for hardware and software.
 - c. The details for hardware include device name, site name, hardware conformance status, model, role, hardware end of vulnerability support for a particular device. Click the device name to view additional details.
 - d. The details for software include device name, site name, software conformance status, model, software version, role, software end of vulnerability support for a particular device. Click the device name to view additional details.
 - e. Use search to filter by attributes such as device, site, hardware conformance status, software conformance status, model, software version, and role.
 - f. Use the gear icon to customize the columns in the table.
6. The Scale page displays all sites summary, scale conformance, and scale metrics.

- a. The All Sites Summary section displays overall scale conformance level, top 5 switches by scalability metric violations, scalability metrics for controller and switches, and total scalability metrics violations.
 - b. Click **View Conformance Criteria** to learn more.
 - c. The Scale Conformance section displays the scale conformance for controller and switch in the last 6 months if the scale reports for previous months are available.
 - d. The All Scale Metrics section displays the scale metrics details for sites and switches. The All Scale Metrics section displays if you choose **All sites** from the drop-down menu.
 - i. The details for sites include site name, type, software version, controller metrics conformance, switch metrics conformance. Click the site name to view additional details.
 - ii. The details for switches include switch name, site name, software version, model, forward scale profile, metrics conformance. Click the switch name to view additional details.
 - iii. Use search to filter by attributes such as site, type, software version.
 - iv. Use the gear icon to customize the columns in the table.
 - e. The Site Level Scale Metrics and Switch Level Scale Metrics displays the scale metrics details for a site and switches associated with the site. These sections are displayed, if you choose one site from the drop-down menu.
 - i. The details for a site include metric, conformance status, and resource usage,
 - ii. The details for switches include switch name, site name, software version, model, forward scale profile, metrics conformance. Click the switch name to view additional details.
7. From the Actions menu, click **Run Report** to run an on-demand report.

Connectivity Analysis

Connectivity Analysis

Connectivity Analysis feature enables you to run an analysis for a flow within the boundaries of a given site. It is a micro-service launched through Nexus Dashboard Insights, used for tracing end-to-end forwarding path for a given flow and narrowing down the offending device on its path.

Connectivity Analysis detects and isolates offending nodes in the network for a given flow and includes the following functionalities:

- Traces all possible forwarding paths for a given flow across source to destination endpoints.
- Identifies the offending device with issue, resulting in the flow drop.
- Helps troubleshoot to narrow down the root cause of the issue, including running forwarding path checks.

The Nexus Dashboard Insights agent is an app running on the Nexus Insights Cloud Connector, which is pre-installed with the Cisco Application Policy Infrastructure Controller (Cisco APIC). The Nexus Dashboard Insights agent gets the path for a specific flow. The job uses the active path returned to go to the next hop running the connectivity analysis job.

The checks performed for Connectivity Analysis include:

- Connectivity Analysis Patch tracer:
 - Topology checks such as overall health, connectivity of leaf switch, spine switch, or remote leaf switch
 - VRF and BD mappings for endpoints
 - Interfaces connectivity such as PC, VPC, SVI, Breakouts, Sublfs
 - Routing tables, EPM, and EPMC tables
 - L3Out information and mapping
 - Adjacency (ARP) tables
 - Tunnel TM information
 - Synthetic routes (COOP) tables on spine switches
- Analyze Active Patch option:
 - Policy tables (ACL)
 - HAL layers
 - ToR glean information
 - ELAM drop codes

Guidelines and Limitations

- At a time, you can submit up to 10 jobs per site.
- At any point of time, you can run only 1 connectivity analysis job per site. You can stop a job in the queue and run another job.

- Connectivity Analysis feature is supported on Cisco APIC release 6.0(2h) and Cisco ACI Switch release 16.0(2h) and later. Cisco Nexus Insights Cloud Connector (NICC) app version 3.0.0.350 is pre-packaged with Cisco APIC release 6.0(2h) and is required for this feature.

Supported Topologies

- Endpoint combinations:
 - EP-EP
 - EP-L3Out
 - L3Out-EP
 - L3Out-L3Out
- Conversation types:
 - L2, L3, L4 (TCP/UDP)
 - V4 and V6 support
 - Transit and Proxy flows
 - Shared Service
- Topologies:
 - Single-Pod and Multi-Pod
 - Remote Leaf-Direct
 - M-Topology (stretched fabric design)
 - vPC
 - 3-tier architectures

Schedule a Connectivity Analysis

Navigate to **Analyze > Analysis Hub > Connectivity Analysis > Create Connectivity Analysis**.

General

1. Enter a **Connectivity Analysis Name** and select a **Site**

Analysis Details

1. Choose **L2** or **L3** routed flow.
2. Enter the mandatory fields and optional fields to configure the job.

Connectivity Analysis Job	Input Fields
L2 bridged flow	Mandatory: Tenant, Source MAC, Destination MAC, BD Name, and VRF. Optional: Source Port, Destination port, Protocol

Connectivity Analysis Job	Input Fields
L3 routed flow	<p>Mandatory: Source Tenant, Source IP, Destination IP, and Source VRF</p> <p>Optional: Destination Tenant, Source Port, Protocol, Destination Port, Protocol, Destination VRF</p>

Select the **Analyze Active Path** check box if you want the Connectivity Analysis to analyze an available active flow to provide additional connectivity information. If you use this option, the analysis may take longer, as a long live flow is needed for the effectiveness of the tool.

You can enable ELAM through 'Analyze Active Path' option.

Summary

This displays a summary of all the data entered to run the analysis job.

The post success screen comes up options to **Create Another Connectivity Analysis** or **View Connectivity Analysis**.

Connectivity Analysis Dashboard

The **Connectivity Analysis** Dashboard displays the list of connectivity analysis jobs along with a graph of the analyses by job and flow status. You can view the analyses based on any selected timeline. By default, the timeline is set to latest.

The Job and Flow Status are color coded. The following table lists the color along with the corresponding status.

Color	Status
Red	Failed
Green	Complete
Yellow	Queued
Blue	In Progress
Grey	Status Unknown

The filter bar allows you to filters the analysis by the following objects:

- Job status
- Flow status
- Name
- SRC IP

- DEST IP
- Analyze Active path
- Creation Time
- End time
- Source MAC
- Destination MAC
- Source port
- Destination port
- Protocol

See [Filtering Information](#) to view the available operators that can be used to filter the information.

The page displays the Connectivity Analysis jobs in a tabular form and are sorted by status. You can customize the view by hiding/showing some of the columns of the table using the gear icon.

The following fields are available for each analysis:

- Job status
- Flow status
- Name
- SRC IP
- DEST IP
- Analyze Active path
- Creation Time
- End time
- Source MAC
- Destination MAC
- Source port
- Destination port
- Protocol

The **Refresh** button allows you to refresh the page to view any new analyses that have run in the background. The **Create Connectivity Analysis** button allows you to create a new analysis. Select any connectivity analysis job in the table to display additional details.

The **Connectivity Analysis** status displays the following information:

- Analysis Results
- Path Summary

Analysis Results

✓ Job Status Completed ✓ Path Status Success

Details

Creation Time	End Time	Run Time	Tenant Name	VRF Name	Destination Tenant Name
Sep 05 2023, 12:02:23.000 AM	Sep 05 2023, 12:05:09.000 AM	1 Minute 19 Seconds	bharat	bharat	-
Destination VRF Name	Source IP	Destination IP	Source Port	Destination Port	Protocol
-			-	-	-

Analyze Active Path
Enabled

Path Summary

© Cisco Systems, Inc.

Analysis Results

Analysis Results shows the Job and Path status of the node. It also displays the following details about the analysis:

- Creation Time
- End Time
- Run Time
- Source IP
- Destination IP
- Source VLAN
- VRF Name
- Source MAC
- Destination MAC
- Source Port
- Destination Port
- Protocol
- Flow Type
- Run Type
- Analyze Active Path - This shows the active path status. Active path is the one where the actual traffic flows through the specific interfaces of all the available interfaces.

Path Summary

The path summary displays a workflow of the path where the traffic flows from the source port to the destination port. The green circle around each node signifies that the nodes are included in the active path. Double-click on a particular node to view the relevant details about it. The +, -, and the **Reset** allows you to zoom in, out and reset the size of the path summary.

The Real Path from topology can be calculated if active flow exists. This is created only when the Source and Destination ports are filled. It is available as a preview graph in grey when the job is in progress. Once completed the path becomes green and the the nodes in the topology graph will be clickable.

Node Details

The following information is available for the nodes:

- General
 - Job ID
 - Source Address
 - Destination Address
 - Source VLAN
 - VRF Name
 - Source MAC
 - Destination MAC
 - Source Port
 - Destination Port
 - Protocol
 - Flow Type
 - Run Type
 - Analyze Active Path
- Details
 1. Paths - Path details for the node. This displays the Ingress Interfaces and the Egress Paths available for the node.
 2. ELAM - ELAM details. If ELAM is triggered, then the report is available. To view the full report, click **View Full Report**. The ELAM report captures the following information:
 - Basic Information
 - Inner L2 Header
 - Inner L3 Header
 - Capture Packet
 - FPX Latch Results
 - Flood/ Multicast or Forwarding to Remote Leaf

- Forwarding Lookup
- Lookup Blocks Information
- Rewrite Information
- iFabric Information
- Lookup Drop
- Other Forwarding Information
- Sideband Information
- Sideband SB Information
- Sideband SB Multicast Information
- SUP-TCAM (ACX)
- Table Lookup Results
- Trigger/Basic Information

The path summary can also be displayed in a tabular form with the menu icon. The following fields are available for each hop of the node:

- Hop Number
- Node name
- Node ID
- Site
- State Validator - This tells whether the current state is a success or not.
- Forwarding - This tells whether the forwarding was successful or not.
- Ingress paths
- Egress paths
- Tunnel Type

Click the **Actions** drop-down menu to:

- Re-run analysis: To re-run the connectivity analysis.
- Run Reverse Flow analysis: To run the reverse flow analysis.
- Show Event log: To display the event log.

Connectivity Analysis Error Scenarios

A Connectivity Analysis job may fail if:

- There is no active flow traffic in the fabric. Ensure that the traffic flows are active when enabling the active path analysis.
- The destination address is not known in the VRF. Ensure that the destination endpoint is configured in VRF.
- Traffic drops in ELAM. This is visible in the path summary.

- The consistency checker fails.

Filtering Information

In some cases, you might be able to filter results to find information more easily.

For example, you might have a situation where there a large number of endpoints under a single leaf switch, but you are only interested in endpoints that have a certain VLAN value.

You could filter the information to show only those specific endpoints in this situation.

Use the following operators for the filter refinement:

Operator	Description
==	With the initial filter type, this operator, and a subsequent value, returns an exact match.
!=	With the initial filter type, this operator, and a subsequent value, returns all that do not have the same value.
contains	With the initial filter type, this operator, and a subsequent value, returns all that contain the value.
!contains	With the initial filter type, this operator, and a subsequent value, returns all that do not contain the value.
<	With the initial filter type, this operator, and a subsequent value, returns a match less than the value.
< =	With the initial filter type, this operator, and a subsequent value, returns a match less than or equal to the value.
>	With the initial filter type, this operator, and a subsequent value, returns a match greater than the value.
> =	With the initial filter type, this operator, and a subsequent value, returns a match greater than or equal to the value.

Delta Analysis

Delta Analysis

Nexus Dashboard Insights performs analysis of sites at regular intervals and the data is collected at an interval depending on the number of nodes.

Number of nodes	Interval
Fewer than 100	2 hours
100 to 400	3 hours
Greater than 400	12 hours

At each interval, Nexus Dashboard Insights captures a snapshot of the controller policies and the fabric run time state, performs analysis, and generates anomalies. The anomalies generated describe the health of the network at that snapshot.

Delta analysis enables you to analyze the difference in the policy, run time state, and the health of the network between two snapshots.

Create Delta Analysis enables you to create a new delta analysis and manage existing analysis. See [Create Delta Analysis](#).

Health Delta

Health Delta analyses the difference in the health of the fabric across the two snapshots.

See [Health Delta](#) for more details.

Policy Delta for ACI

Policy Delta analyzes the differences in the policy between the two snapshots and provides a correlated view of what has changed in the ACI Fabric.

See [Policy Delta](#) for more details.

Guidelines and Limitations

- The Delta Analysis functionality currently supports the local authentication domain only.
- While you are currently allowed to create more than one Delta Analyses at any given time, we recommend that you do not queue more than one Delta Analysis at any given time. In addition, we recommend that you wait for some time (approximately 10 minutes) between creating new analyses to avoid the risk of adversely impacting the run time of the concurrent online site analysis.

The interdependency arises because the Delta Analysis results in an increased load on the database. Sustained high-database load from multiple back-to-back Delta Analyses may affect the run-time of the online analysis.

- The **APIC Configuration Export Policy** must be of the same format (XML/JSON) for both the snapshots.
- The policy delta will not be performed if there are any APIC configuration export policy collection errors.

Create Delta Analysis

For ACI Assurance Group users, APIC admin writePriv privileges allow information collection on the APIC host and leaf switches. You must have APIC admin writePriv privileges to configure the **APIC Configuration Export Policy**.

Choose **Analyze > Analysis Hub > Delta Analysis > Create Delta Analysis**.

1. In the **Delta Analysis Name** field, enter the name. The name must be unique across all the analyses.
2. Click **Site** to select the site.
3. Click **Select Earlier Snapshot** and choose the first snapshot for the delta analysis. Click **Apply**.
4. Click **Select Later Snapshot** and choose the second snapshot for the delta analysis. Click **Apply**.



The two snapshots selected for the delta analysis must belong to the same site.

5. View the Summary of the Delta Analysis created in **Summary**.
6. Click **Save**. The status of the delta analysis is displayed in the **Delta Analysis** table. Post completion allows you to **View Delta Analysis** or **Create another Delta Analysis**.

You can perform one delta analysis at a time. To perform another delta analysis, you must stop the current delta analysis and then start the another delta analysis.

1. (Optional) From the Status column, select an In Progress or Scheduled analysis and click **STOP** in the "..." option to stop the delta analysis.
2. The **Delete** in the "..." allows you to delete the analysis created.



If there are any errors in the creation of a delta rule, it will be displayed on the summary of the rule creation as a banner.

View Delta Analysis

The page displays the analysis in a tabular form. The analysis are sorted by status. The **Create Delta Analysis** button lets you create a new delta analysis.

The status of analysis can be either **Aborted, Pending, Scheduled, Stopped, Stopping, Success, Failed, Partially Failed, Queued, Completed** or **In progress**.

The filter bar allows you to filters the analysis by the following factors:

- Name
- Status

- Site
- Submitter ID

The delta analysis dashboard displays the general details of the analysis along with the health and policy delta.

- To view the results of health delta analysis, see [View Health Delta Analysis](#).
- To view the results of policy delta analysis, see [View Policy Delta Analysis](#).

View Health Delta Analysis

Health Delta analyses the difference in the health of the fabric across the two snapshots. The results are displayed in the following areas:

- **Anomaly Count:** Displays the difference in anomaly count per severity across the snapshots. If you click on the difference that shows, the **All Anomalies** table gets filtered accordingly. The first count represents the anomalies found only in the earlier snapshot. The second count represents the anomalies common in both the snapshots. The third count represents the anomalies found only in the later snapshot.
- **Health Delta by Resources:** Displays the count of resources by type that have seen a change in their health. You can also view the resources whose health has changed by checking the **View Changed** checkbox. The gear icon allows you to customize the columns as per your view. The filter bar helps to filter the resources in the table by 'Resource'.

The table displays count delta and health delta. Count delta includes both healthy and unhealthy resources. Healthy resources prob won't have any anomalies associated with it if filtered Health delta shows only unhealthy resources and should return anomalies if filtered by

- **All Anomalies:** The **Grouped** view displays the delta status for grouped anomalies across the snapshots. The **Ungrouped** view displays the delta status for each anomaly across the snapshots.

The Anomalies can be listed for the following kinds of snapshots:

- Earlier Snapshot
- Later Snapshot
- Earlier Snapshot Only
- Later Snapshot Only
- Both Snapshots

The anomalies are displayed in a tabular form with the following fields:

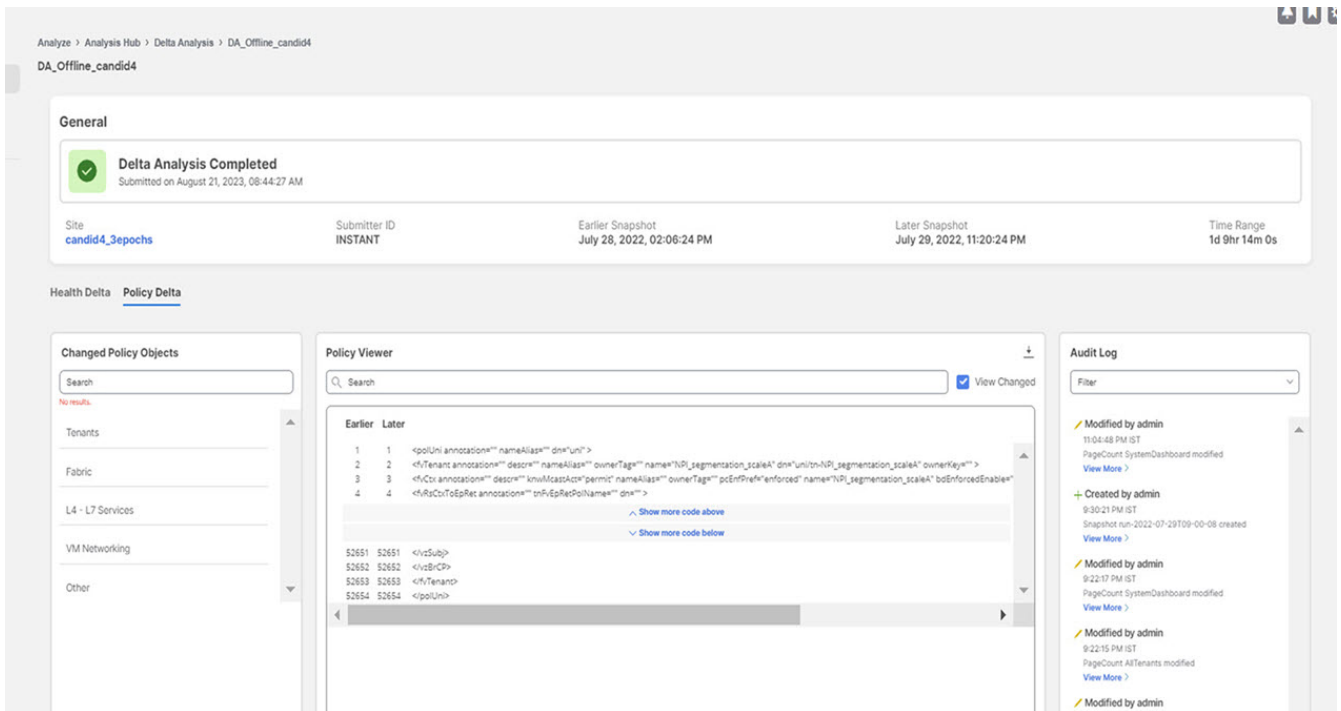
- Title
- Anomaly Level
- Category
- Count

The gear icon allows you to customize the columns as per your view.

You can filter the results based on the following attributes:

- Anomaly Level
- App Profile DN
- BD DN
- Title
- Contract DN
- EPGs
- External Routes
- Interfaces
- Internal Subnets
- L3Out DN
- Leaf DN
- Tenant DN
- Endpoints
- VRF DN

View Policy Delta Analysis



Click **Policy Delta** to view the policy changes across the two snapshots. Policy Delta includes three sections: Changed Policy Object, Policy Viewer, and Audit Log.

1. The **Changed Policy Object** panel, displays the changed policy object tree across the two snapshots. The corresponding changes in the Policy Viewer and Audit Log panels are highlighted. Use the **Search** bar to perform a DN search.
 - a. Drill down on a particular object to view the object types that have changed. The number indicates the number of changes to the object.
 - b. Select the changed object type to view the anomalies that have changed.
 - c. Click DN link to access the affected object type in APIC.
 - d. Click **Show Changes** to view the changes in the Policy Viewer and Audit Log panels.
2. The **Policy Viewer** panel displays the policy configuration across the earlier and later snapshots. It also helps view the added, modified, and deleted policy configurations between the two snapshots and view the context around the modified areas in the policy delta.
 - a. Use the color coding to visualize the added, deleted, modified, and unchanged content across the two policies.
 - b. Click **Show More Code Above** or **Show More Code Below** to display more content.
 - c. Click the download icon to export the policy configuration for the earlier snapshots policy and later snapshots policy.
 - d. Enter a value in the **Search** bar to perform a text search in added, modified, deleted, and unchanged areas in the policy delta.
3. The **Audit Log** panel then displays all the audit logs that were created between the two snapshots. Cisco Nexus Dashboard Insights collects audit logs from APIC and computes the difference in the audit logs between the two snapshots.

A correlated view of what has change in the data center is displayed in the **Audit Log** panel. When you select a particular object in the **Changed Policy Objects** panel, the relevant difference is highlighted in the **Policy Viewer** panel and the relevant audit log is highlighted in the **Audit Log**

panel. APIC audit logs are records of user-initiated events such as logins and logouts or configuration changes that are required to be auditable. For every snapshot, the audit log history is limited to last 24 hrs.

- a. Use the **Filter** bar to filter by DN, User ID, or Any.
- b. Click **View More** on an audit log entry to view when the changes were made and who made the changes. The timestamp on the audit log entry corresponds to the the timestamp on the APIC audit log.
- c. Click Audit Log entry to access the affected object type in APIC.

Log Collector

Log Collector

The Log Collector feature enables you to collect and upload the logs for the devices in your network to Cisco Intersight Cloud. It also enables Cisco TAC to trigger on-demand collection of logs for devices on the site and pulls the logs from Cisco Intersight Cloud.

The Log Collector has two modes:

- User initiated - The user collects the logs for devices on the site and then uploads the collected logs to Cisco Intersight Cloud after the log collection job is completed. You can automatically upload the log files to Cisco Intersight Cloud after the log collection job is completed.
- TAC initiated - Cisco TAC triggers on-demand collection of logs for specified devices and pulls the logs from Cisco Intersight Cloud.

Device Connectivity Notifier for TAC Initiated Collector

Nexus Dashboard Insights uses the device connectivity issue notifier on Cisco Nexus Dashboard to communicate with the devices. The notifier checks for TAC triggered on-demand collection of logs. In case the fabric is not configured properly to communicate with the device, Nexus Dashboard Insights notifies the following:

- The device is not configured for node interaction.
- You can not run a Log Collector job on the device.
- Nexus Dashboard Insights cannot connect to the device.

If the node interaction is not healthy on the device, you cannot select the device for Log Collector to collect logs. In the GUI, the device is greyed out.

Log Collector Dashboard

Navigate to **Analyze > Analysis Hub > Log Collector**.

The **Log Collector** Dashboard displays a graph of Logs by Job status for a particular site and displays the latest log collections. The filter bar allows you to filters the logs by Status, Name, Node, start time, and end time.

Use the following operators for the filter refinement:

Operator	Description
==	With the initial filter type, this operator, and a subsequent value, returns an exact match.
!=	With the initial filter type, this operator, and a subsequent value, returns all that do not have the same value.

Operator	Description
contains	With the initial filter type, this operator, and a subsequent value, returns all that contain the value.
!contains	With the initial filter type, this operator, and a subsequent value, returns all that do not contain the value.

The page also displays the log collection jobs in a tabular format. The jobs are sorted by status. Select the log collection job in the table to view additional details.

General

This displays the status of the job along with a graph showing the number of devices by status.

Details

The following information is listed:

- Creation Time
- End Time
- Nodes
- Job ID

Selected Nodes

This displays the list of nodes in a tabular form along with the status of each job and the upload status for the files uploaded.



Upload All Files allows you to upload all the files.

... allows you to Download each file separately.

TAC Initiated Log Collector

The TAC initiated log collector enables Cisco TAC to trigger on-demand collection of logs for specified user devices in the Cisco Intersight Cloud to the Device Connector.

When the TAC assist job is complete, the new job appears in the **Log Collector** table. Select the log collection job in the table to display additional details. **Log Collection** status displays information such as status, general information, and node details.

Upload logs to Cisco Intersight Cloud



Ensure that Nexus Dashboard Insights is connected to Cisco Intersight Cloud and Cisco Intersight Device Connector.

Choose **Analyze > Analyze Hub > Log Collector > New Log Collector**.



1. Enter the name.
2. Click **Select Site** to select a site.
3. (Optional) Check **Auto Upload Log Files** to automatically upload the log files to Cisco Intersight Cloud after the log collection job is completed.
4. Click **Next**.
5. Click **Add Nodes** and then select the nodes from the **Select Nodes** menu.
6. Click **Add**. The nodes are displayed in the **Select Nodes** table.



For a physical Nexus Dashboard, you can select up to 10 nodes.

1. Click **Start Collection** to initiate the log collection process.

When the job is complete, the new job appears in the **Log Collector** table.

2. Click the job in the table to display additional job details.
3. Click the  icon to display **Log Collection** status.
4. Select the node and click  icon.
5. Click **Upload File to TAC Assist** to upload a single file for the selected node manually.
6. Click **Upload** to upload all the log files generated for the selected node manually.

The status of the upload is displayed in the **Selected Nodes** table.

Guidelines and Limitations

- If the upload logs fails for some of the nodes and succeeds for the rest of the nodes, then in the **Selected Nodes** table, the status is displayed as Completed.
- If the upload is in progress, the status is displayed as Queued.
- If the collection fails for some of the nodes, then the collection will continue for other nodes. After the collection is completed, the upload will start. In the **Selected Nodes** table, the combined status is displayed in the Status column.
- If the collection succeeds for some of the nodes, but the upload fails, then in the **Selected Nodes** table, the status is displayed as Failed.
- **Auto Upload Log Files** can be performed only on one node at a time.

Compliance

Compliance

The **Compliance** feature enables the user to accomplish compliance results. There are two Compliance Types - **Communication** and **Configuration** Compliance.

- **Communication Compliance** consists of the following Compliance Rule Types:
 - Service Level Agreement (SLA) Compliance: You can set up rules for entities that must talk with other entities. You can use the Compliance feature to set up regulatory compliance rules.
 - Traffic Restriction Compliance: You can specify restrictions on protocols and ports for communication between objects.
 - Segmentation Compliance: You can establish walled areas around a set of entities that must not communicate with other entities.
- **Configuration Compliance** - You can perform a configuration compliance check against a specified configuration.

The four types of Configuration compliance that can be performed are:

- Snapshot Settings
- Manual Configuration
- Template based Compliance
- Import Configuration

In the UI, you specify your compliance rules and Cisco Nexus Dashboard Insights will verify in the subsequent snapshots, whether the compliance rules are satisfied by the policy that is configured on Cisco APIC.

If satisfied, Nexus Dashboard Insights raises one anomaly per snapshot stating that the compliance rule is satisfied. For example, if there are two rules associated with the snapshot, Nexus Dashboard Insights raises two anomalies.

An assurance group runs a compliance analysis on a snapshot at an interval depending on the number of nodes.

Number of nodes	Interval
Fewer than 100	2 hours
100 to 400	3 hours
Greater than 400	12 hours

The following examples provide with information about the compliance **include** and **exclude** rules:

- Contains EPGs in Tenants with names that start with “a” or ending with “z” : EPGs in Tenants such as “abz” that satisfy both criteria are included only once.
- Contains EPGs in Tenants with names that start with “a” and are also in VRFs where the Tenant is “xyz” and the VRF name contains “c” : When an EPG under Tenant “abc” that is in a VRF with DN

uni/tn-xyz/ctx-abcde is selected, verify that both the Tenant and the VRF criteria match. An EPG under Tenant “abc” that is in a VRF with DN uni/tn-xyz1/ctx-abcde is not selected because the VRF Tenant does not match.

- Contains all EPGs under Tenants that begin with “a” except those that contain “d” : An EPG under Tenant “abc” is selected. An EPG under Tenant “abcd” is not selected.
- Contains all EPGs under Tenants that begin with “a” except those EPGs that are also in the VRF with DN uni/tn-rrr/ctx-sss : An EPG under Tenant “abc” that is in a VRF with DN uni/tn-rrr/ctx-sss is selected because the VRF Tenant matches.

Compliance is supported in the following Cisco APIC releases:

- 3.2(x) release
- 4.0(x) release
- 4.1(x) release
- 4.2(x) release
- 5.0(x) release
- 5.1(x) release
- 5.2(x) release
- 5.3(x) release
- 6.0(x) release

Guidelines and Limitations

- A single compliance can be associated with multiple sites. However, traffic selectors and object selectors that are created for one compliance rule cannot be reused by another compliance rule. New selectors must be created each time you want to add a selector to a new compliance rule.
- In the **Compliance Requirement Type** field, the **Configuration** and **Communication** tabs enable and enforce the configuration to meet best practices and business requirements that will be met only if you choose to run your changes through Cisco Nexus Dashboard Insights before you apply those changes on the controller. Otherwise, Cisco Nexus Dashboard Insights will not enforce the compliance but will report it as a violation.

Configuration enables and enforces the configuration to meet best practices and business requirements.

Communication enables communication or isolation between network objects that meet business and regulatory purposes.

- Compliance Rules are created at the site level, and the sites that you choose must be part of that group.
- Compliance action rules are available within the sites where you define the compliance rules and actions.
- You can have a maximum of 30 active Communication Compliance rules and 600 active Configuration Compliance rules per site. If you exceed this limit, you cannot add more requirements in the **Manage Compliance** area.
- When a compliance job is in progress for one or more sites, do not start a bug scan for those

sites.

Create Compliance Communication Rule

Navigate to **Manage > Rules > Compliance Rules > Create Compliance Rule**.



An alternate way to get to the rule creation page, is to click the Create Compliance Rule button in **Compliance** in **Analysis Hub**. This will take you to rule creation.

Once you're on the rule creation page, follow these steps:

Basic Information

1. Provide the name and description for your rule. You can choose to Enable or Disable the rule.



When you create a compliance rule, you can add a custom description, which appears in the compliance violation anomaly.

If the rule is in Enabled state, the rule will be used to generate the Compliance Report, the next time it gets generated. If the rule is in Disabled state, it will not be used.



The state of any rule can be changed after they have been created. Go to **Configure > Rules** and enable or disable the specific rule from the Rule State column in the table.

Rule Naming Criteria:

- o Name should be a minimum of three characters
 - o Name should not include special characters
 - o Name should be unique.
 - o No two rules can have the same name.
2. Select the sites you would like to apply the rule to. You can pick one, or many, or all sites.

Settings

1. In the **Compliance Rule Type** field, choose **Communication**.
2. Under **Criteria**, for the **Communication Type** field, choose the appropriate communication type.

The options are **Must Talk To**, **Must Not Talk To**, **May Talk To**. For more details about which Compliance Rule Type is configured based upon your selections, see the table that follows this procedure.

3. In the **Object Type** fields and the **Traffic Selector** area, choose the appropriate objects and traffic selector.

The Communication Types are applied between two different object groups.

4. Select the appropriate criteria for both groups.
5. After you define the criteria in the **Add Criteria** area, click the **View Selected Objects** link, and verify that the selected objects are appropriate. Based upon your selections of Communication Type and Traffic Selector Rules, the Compliance Rule Type that you defined will be displayed.
6. After you complete defining the objects, criteria, traffic restrictions as appropriate for your site/s, click **Save Rule** to complete the configuration.

The following table displays which Compliance Rule Type is configured based upon your selections of Communication Type and Traffic Selector Rules.



Additional descriptions about the Communication Types follow the table.

Table 1. Communication Type and Traffic Selector Rules Selections and the Resultant Compliance Rule Type

Communication Type	Select a Traffic Selector Rule?
Objects You Can Select	Compliance Requirement Type
Must Talk To	Mandatory to select
EPG	Service Level Agreement (SLA)
Must Not Talk To	Not mandatory to select
<ul style="list-style-type: none"> • EPG • Tenant 	<ul style="list-style-type: none"> • If you select a Traffic Selector Rule, the Compliance Rule is Traffic Restriction • If you do not select a Traffic Selector Rule, the Compliance Rule is Segmentation
May Talk To	Mandatory to select
EPG	Traffic Restriction

Must Talk To: This allows you to configure objects where *selector A* **must talk to** objects selected by *selector B* under defined traffic restriction rules.

Must Not Talk To: Choose this configuration if your intention is that an object selected by object *selector A* **must not talk to** objects selected by object *selector B* using a defined type of traffic. The traffic restriction rule is optional in this configuration. Two different types of communication compliances can be configured using this option:

- Traffic Restriction compliance: You can specify a traffic selector rule that objects selected by *selector A* **must not talk to** objects selected by *selector B*, using a selected type of traffic that uses traffic restriction rules. This communication is restricted
- Segmentation compliance: By not defining a traffic selector rule, you can configure segmentation compliance where objects in *selector A* **cannot talk to** objects in *selector B* using any type of traffic. In this case, no traffic restriction rules are defined by you.

May Talk To: This allows you to create a traffic restriction compliance. Objects selected by *selector A* **may talk to** objects selected by *selector B* using only a specific type of traffic using traffic restriction

rules. As an example, EPG A and EPG B are connected, but they can talk to each other using only the specific traffic types that are defined in the Cisco APIC contract using filters. As a Nexus Dashboard Insights user, to verify that EPG A can talk to EPG B using the traffic type TCP IP, configure the traffic restriction rule EPG A **May Talk To** EPG B using TCP IP.

Matching Criteria

Select any object type and the corresponding matching criteria object. See [Matching Criteria](#) for the available object types.

To understand how the various matching criteria objects can be defined, see [Matching Criteria](#).

The various traffic selector rules available are:

Ether Type	Protocol Type
ARP	-
FCOE	-
IP	<ul style="list-style-type: none"> ▪ All ▪ EGP ▪ EIGRP ▪ ICMP ▪ ICMPV6 ▪ IGMP ▪ IGP ▪ L2TP ▪ OJPFIGP ▪ PIM ▪ TCP ▪ UDP
MAC_SECURITY	-
MPLS_UNICAST	-
TRILL	-



You can view/edit Direction based traffic settings from the **Direction settings** column.

Summary

You can view the entire overview of the rule you want to create and click **Save** when you are ready to create the rule. When the rule is saved, you see the post success screen. You can choose to **View compliance rules**, **View Compliance**, or **Create another Compliance rule** from this page.

Create Compliance Rule with Snapshot Selection



This is similar to the Configuration Compliance Check method but you also select a snapshot. With this method, you can make sure that certain attributes of objects are not changed when going from one snapshot to another snapshot.

1. Under **Compliance Rule Type**, choose **Configuration**.
2. In the **Base Configuration Settings** field, choose **Snapshot Settings**.
3. In the **Time of Snapshot** field, choose the desired snapshot time, and click **Apply**.
4. In **New Rule**, click **Save**. Cisco Nexus Dashboard Insights starts performing a check.

To download the snapshot, click the **Download** link from **Settings**.

Create Import Configuration Compliance

You can perform an import configuration against a specified configuration. You specify a configuration file or snapshot, and Cisco Nexus Dashboard Insights continuously checks against it and enables you to identify changes for the objects and configurable attributes defined in Cisco APIC. If the configuration deviates from the specified configuration, then violations are raised. For each violation, there will be a separate violation anomaly displayed. Additionally, a single anomaly will be raised that includes every variable for every object of the Tenant that is not a violation.

Once you've selected the **Compliance Rule Type** as *Configuration* follow this step:

1. In the **Base Configuration Settings** field, choose **Import Configuration** and drag and drop your file into the provided field to upload. Click **Save**.

You can check the box to allow addition of new configuration objects. This will raise a violation for every new object which is missing in the uploaded configuration file.



You cannot edit the configuration rules when you upload a JSON/XML file. In such a case, after uploading a file, you can view or download it by navigating from **Actions**.

Template Based Compliance

With Template Based Compliance, you have the flexibility to select objects based on any attributes and provide different types of matching criteria that are not supported when you configure other compliance tasks.

Template-based compliance allows you to configure a template and specify types of queries to select objects and attributes that enforce specific conditions when enabled. The Template Query Language enables you to select any configurable object and define what attributes to apply to the compliance. **SELECT** allows you to choose a certain subset of objects that you want to specify, and **MATCH** defines the Compliance rule.

With other types of Compliance configurations releases you can upload a JSON/XML file and all the attributes in the file will be matched as is. Alternatively, you can also select a few specific objects based on name matches, and you can configure select attributes supported for those specific objects. This allows you to search for existing or future objects matching the names that are checked for compliance for the specified parameters.

Verified Scalability Limits for Template Based Compliance

- Number of Template rules are 5 for APIC with total configurable objects of 150,000.
- Each template selects 15,000 objects on an average.
- Number of tenants per template is 30 tenants, with each tenant selecting 500 objects on an average.
- You may create more than 5 templates (the upper limit is 30 total rules), if the total objects selected by all the templates are less than 5*15,000 and the total configurations in APIC are < 150,000 objects.
- You can have a maximum of 30 active Communication Compliance rules and 600 active Configuration Compliance rules per site.

Template Guidelines

Follow these guidelines when defining a template:

- The JSON file format is supported for the template query. The XML format is *NOT* supported.
- The template follows the same structure as used in APIC files. It has objects, attributes, and children.
- The template file size that you upload can be up to 15 MB including whitespaces. Pretty JSON files will have whitespaces to support indentation. To reduce the file size, you can remove whitespaces and upload the file.

Template Syntax Guidelines

Follow these syntax examples and guidelines:

Template Syntax for SELECT

SELECT allows the user to choose a subset of objects based on the criteria defined using the KEYWORDS. In the following example, **SELECT** is followed by the attribute name which is an attribute

of the object. You must use one of the following keywords:

- STARTS_WITH
- ENDS_WITH
- EXACT
- OR
- REGEX

Syntax:

SELECT(<attribute_name>): KEY_WORD(<value>)

Attribute_name any attribute of the object.

REGEX(<value>) - where the value must follow the standard regex expression syntax
"SELECT(name)": "REGEX(Ctrct_[1-3])"

For more details about keyword regular expressions, see [Summary of Regular-Expressions Constructs](#).

The following is a syntax example:

```
{
  "fvAEPg":
  {
    "attributes":
    {
      "dn": " EXACT(uni/tn-aepg_vzanycons_imd_ctx_pass_7/ap-CTX1_AP1/epg-
CTX1_BD1_AP1_EPG7)",
      "MATCH(isAttrBasedEPg)": " EXACT(no)"
    }
  }
}
```

If **SELECT** is not specified for an attribute, then **rn** and **dn** will be considered as **SELECT** by default.

Template Syntax for MATCH

The **MATCH** criteria is used to define the Compliance rule. These compliance rules will be applied on objects that are selected using the **SELECT** criteria. The keyword to match the attribute is MATCH, and the match is applied to all objects of the specified type. Specify the attributes and values that you want to match. The attributes are matched to all the objects of the specified type.

The values are not required to be exact matches. They can be as follows:

- STARTS_WITH
- ENDS_WITH
- EXACT

- OR
- REGEX

MATCH(<attribute_name>): KEY_WORD(<value>)

The following is a syntax example where you select all **vzBrCP** contract objects where the name is **Ctrct_1** or **Ctrct_2** or **Ctrct_3**. Then you define that the scope is context.

```
"vzBrCP": {
  "attributes": {
    "SELECT(name)": "REGEX(Ctrct_[1-3])",
    "MATCH(scope)": "EXACT(context)"
  }
}
```

The following is a syntax example where if the KEY_WORD is not defined, the default behavior is **EXACT**. When you use MATCH(dn) and MATCH(rn), they are defined as match criteria.



If an attribute (other than **dn** and **rn**) does not have **MATCH** or **SELECT** specified, it will be considered as **MATCH** by default.

```
{
  "fvAEPg": {
    "attributes": {
      "SELECT(dn)": "uni/tn-aepg_vzanycons_imd_ctx_pass_7/ap-CTX1_AP1/epg-
      CTX1_BD1_AP1_EPG7",
      "MATCH(isAttrBasedEPg)": "EXACT(no)",
      "prio": "OR(unspecified, prio1)"
    }
  }
}
```

In the above example, by default, "prio" will be a **MATCH**.

Template Syntax for {}

The following is a syntax example of a generic template where the **KEY_WORD** is {}. You can use this template to customize your requirements, select attributes, regular expressions.

The **KEY_WORD** values can be as follows:

- STARTS_WITH
- ENDS_WITH
- EXACT
- OR

• REGEX

```
{
  "<MO type>": {
    "attributes": {
      "SELECT(<attribute>)": " KEY_WORD(<expression>)",
      "MATCH(<attribute>)": " KEY_WORD (<value>)"
    },
    "children": [
      {
        "<MO type>": {
          "attributes": {
            "SELECT(<attribute>)": " KEY_WORD (<value>)",
            "MATCH(<attribute>)": " KEY_WORD (<value>)"
          },
          "children": [
            {
              "<MO type>": {
                "attributes": {
                  "SELECT((<attribute>)": " KEY_WORD (<value>)",
                  "MATCH(<attribute>)": " KEY_WORD (<value>,<value>)"
                }
              }
            }
          ]
        }
      }
    ]
  }
}
```

EXAMPLE OF TEMPLATE BASED CONFIGURATION COMPLIANCE

The following is an example of a Template Based Configuration Compliance. In this example, choose all the contracts where **name** starts with **Ctrct_(1-3)**. Then, match **scope** which must be **context**. For **children** of the subjects of those contracts, select **name** as any (wildcard) and **nameAlias** must be ABC.

```
{
  "vzBrCP": {
    "attributes": {
      "SELECT(name)": " REGEX(Ctrct_[1-3])",
      "MATCH(scope)": " EXACT(context)"
    }
  }
}
```


- In a template, defining **attributes** is mandatory because the Compliance is applied on the attribute.
- In a template, defining **children** is optional. If children are defined in the query, the selection is applied to the real children of the selected objects.
- In a template, you can include the same object type only once per child array. This prevents the possibility of creating requirements that will result in conflicting compliance rules that result in violation anomalies.

INVALID EXAMPLE

In the following *invalid* example, if there is a BD named **ABCXYZ**, it will be selected by both the child object templates snippets for **fvBD**, and one of them will be a violation because **type** can either be **regular** or **fc**.

+

```
{
  "fvTenant": {
    "attributes": {
      "SELECT(name)": " EXACT(tenantABC)"
    },
    "children": [
      {
        "fvBD": {
          "attributes": {
            "MATCH(type)": " EXACT(regular)",
            "SELECT(name)": " REGEX(. *ABC.*)"
          }
        }
      },
      {
        "fvBD": {
          "attributes": {
            "MATCH(type)": " EXACT(fc)",
            "SELECT(name)": " REGEX(. *XYZ.*)"
          }
        }
      }
    ]
  }
}
```

Configure Template Based Compliance

1. In the **Base Configuration Settings** field, choose **Template Based Compliance**.

2. In the **Choose a file or drag and drop to upload** area, upload your template based file. After the file upload is complete, you can click the View icon to review the contents of the file that you uploaded.



A JSON file is currently supported. XML file is not supported. The template file size that you upload can be up to 15 MB. The view feature will not be available if the file size is greater than 5 MB. If the file size is greater than 5 MB, you can download the file and view the contents.

3. Click **Save**.

If the **Status** is **Enabled**, the compliance rule is ready to be consumed in the next snapshot. If the status is **Disabled**, you can click the Actions menu for the row and click **Edit** to open **Edit Compliance Rule**. In the **State** field, change the state to **Enabled** and click **Save**.

Use a Template to Configure Object Selectors for Naming Compliance

When you use the [Create Compliance Communication Rule](#) task to configure compliance, only a few specific object selectors are supported (such as BD, EPG, VRF). By using a template, you can configure *any* object for a Compliance Communication Rule.

An object can be any managed object from APIC, and its selection is based on the Distinguished Name of the object. If you prefer to have a different attribute as the selection criteria, you can use any valid attribute of that object. The **SELECT** criteria can be defined using one of the following keywords.

- STARTS_WITH
- ENDS_WITH
- EXACT
- OR
- REGEX

The following is a syntax example:

```
{
  "<object>" :{
    "attributes" :{
      "SELECT(dn)" : "<KEY_WORD>(<value>)",
      "MATCH(nameAlias/name)" : "<KEY_WORD>(<value>)"
    }
  }
}
```

For Naming Compliance, the Compliance rules are on the name and nameAlias fields that are indicated by **MATCH**. The **MATCH** criteria can be again defined using the keywords provided earlier in this section.

Use the following example template to configure a Naming Compliance to match selected objects to **name** or **nameAlias**:

```
{
  "vzSubj" :{
    "attributes" :{
      "SELECT(dn)" : " EXACT(subj1)",
      "MATCH(nameAlias)" : " STARTS_WITH(ABC)"
    }
  }
}
```



In the above template, you can use any object instead of "vzSubj", and you can use any attribute instead of "dn".

As the attribute **dn** is always considered as **SELECT** by default and any other attribute is always considered as **MATCH**, the above template can be simplified as displayed in the example below. Additionally, if the keyword is not defined, the default behavior is **EXACT**.

```
{
  "vzSubj" :{
    "attributes" :{
      "dn" : " subj1" " nameAlias" : " STARTS_WITH(ABC)"
    }
  }
}
```



In the above template, you can use any object instead of "vzSubj", and you can use any attribute instead of "dn".

For the procedure to configure Object Selectors for Naming Compliance using the above template, see [Configure Template Based Compliance](#).

Use a Template to Configure Object Selectors Based on Tags and Annotations

As an APIC user, you can create tags on managed objects (MOs) that result in creating child objects of type **tagInst** or **tagAnnotation** (based on which APIC version is in use).

Therefore, if you select objects based on a tag created in APIC, you can follow the templates provided in this section to configure object selectors on tags and annotations.

The following is an example that displays the child object as type **tagInst**:

```
{
  "<object>" :{
    "attributes" :{
```

```

    "MATCH(<attribute_name>)" :<KEY_WORD(<value>)"
  },
  "children" :[
    {
      "<tagInst>" :{
        "attributes" :{
          "SELECT(<attribute_name>)" :<KEY_WORD(<value>)"
        }
      }
    }
  ]
}

```

The following is an example that displays the child object as type **tagAnnotation**:

```

{
  "<object>" :{
    "attributes" :{
      "MATCH(<attribute_name>)" :<KEY_WORD(<value>)"
    },
    "children" :[
      {
        "<tagAnnotation>" :{
          "attributes" :{
            "SELECT(<key or value>)" :<KEY_WORD(<value>)"
          }
        }
      }
    ]
  }
}

```

An object can be any valid APIC object with **tagAnnotation** or **tagInst** as a child. Object selection is defined in the **tagInst** or **tagAnnotation** object using **SELECT** on the name in the case of **tagInst**, and **key or value** in the case of **tagAnnotation**. The selection criteria can be any of the following keywords:

- STARTS_WITH
- ENDS_WITH
- EXACT
- OR
- REGEX

Compliance rules are defined at the parent object level using **MATCH** and the criteria can be defined using any **KEY_WORD**. **tagInst** or **tagAnnotation** do not participate in compliance rules as they only provide the selection criteria.

The following is an example template where you **SELECT** all the fvBDs where the tag is “BDs_in_cisco”, and those BDs must have name as **BD** or **app1BD**.

```
{
  "fvBD" :{
    "attributes" :{
      "MATCH(name)" : "OR(BD, app1BD)"
    },
    "children" :[
      {
        "tagInst" :{
          "attributes" :{
            "SELECT(name)" : "EXACT(BDs_in_cisco)"
          }
        }
      }
    ]
  }
}
```

For the procedure to configure object selectors based on Tags and Annotations using a template, see [Configure Template Based Compliance](#).



When using the steps to [Configure Template Based Compliance](#), to configure object selectors for tags and annotations, you must perform an additional step. Before you click **Save**, in **Create New Rule**, you must check the checkbox for the field **Enable Object Selection Based on tagAnnotation/tagInst**. Therefore, if any object has a tag annotation or tagInst, the parent based on the selection criteria in these two objects will be selected.

Create Compliance Rule with Manual Configuration

You can configure this for certain objects such as BD, VRF, EPG, Contract, Subject, and Filter. All objects types are not supported.

Once you're on rule creation, follow these steps:

Basic Information

1. Provide the name and description for your rule. You can choose to Enable or Disable state.



When you create a compliance rule, you can add a custom description, which appears in the compliance violation anomaly.

If the rule is in Enabled state, the rule will be used to generate the Compliance Report, the next time it gets generated. If the rule is in Disabled state, it will not be used.



The state of any rule can be changed after they have been created. Go to **Configure > Rules** and enable or disable the specific rule from the Rule State column in the table.

Rule Naming Criteria:

- o Name should be a minimum of three characters
 - o Name should not include special characters
 - o Name should be unique.
 - o No two rules can have the same name.
2. Select the sites you would like to apply the rule to. You can pick one, or many, or all sites.

Settings

1. In the **Compliance Rule Type** field, choose **Configuration**
2. In the **Base Configuration Settings** field, choose **Manual Configuration**.
3. Under Object Selection, select the **Object Type** and add the criteria as appropriate. You can also view the selected objects with the 'View Selected Objects' button.
4. Click **Save Rule**. Cisco Nexus Dashboard Insights starts performing a check based on the Naming compliance requirements that you specified.

Matching Criteria

Select any object type and the corresponding matching criteria object. See [Matching Criteria](#) for the available object types.

To understand how the various matching criteria objects can be defined, see [Matching Criteria](#).

Configuration Compliance Rules

Add the rules for the matching criteria selected above here. Click 'Add Rule' and select the Attribute, Operator and Value for the rule.



The name and name alias attribute requirement has an additional option to select Matches Regular Expression.

Under attribute requirement you can set the requirement according to the objects selected.

Object	Associated Attributes
EPG	<p>The associated attributes are:</p> <ul style="list-style-type: none"> • Preferred Group Member- The preferred group member can be equal to/not equal to either <i>Include</i> or <i>Exclude</i>. • Infra EPG Isolation- The Infra EPG Isolation can be equal to/not equal to Unenforced/Enforced. • QoS Class- The QoS Class can be equal to/not equal to Unspecified/Level 1/Level 2/Level 3.
VRF	<p>The associated attributes are:</p> <ul style="list-style-type: none"> • Enforcement Preference- The enforcement preference can be set to equal to/not equal to Unenforced/Enforced. • Enforcement Direction - The enforcement direction can be set to equal to/not equal to Ingress/Egress. • Preferred Group - The preferred group can be set to equal to/not equal to Disabled/Enabled. • BD Enforcement - The BD enforcement can be set to equal to/not equal to Yes/No.

Object	Associated Attributes
Bridge Domain (BD)	<p>The attributes are:</p> <ul style="list-style-type: none"> • BD Type - The BD type can be equal to/not equal to regular/FC. The default value is set as equal to regular. • L2 Unknown Unicast - This can be equal to/not equal to Flood/Hardware Proxy. • L3 Unknown Multicast Flooding - This can be equal to/not equal to Flood/ Optimized Flood. • BD Multi Destination Flooding - This can be equal to/not equal to Flood in Encapsulation/Drop/Flood in BD. • PIM - This can be equal to/not equal to Enabled/Disabled. • ARP Flooding - This can be set to equal to/not equal to Yes/No. • Limit IP Learning to Subnet - This can be set to equal to/not equal to Yes/No. • Unicast Routing - This can be set to equal to/not equal to Yes/No. • Subnets - This can be set to All/ None/ At least one to Shared/ Private/ Public.



For BDs in context to VRFs, an extra requirement is needed. The EPG association requirement is to be added which requires an EPG association count. This can be **equal to/at least/at most**. However you can choose to add either the EPG Association Requirement or the Name and Attribute Requirement for BD. You cannot have all the attributes selected.

Summary

You can view the entire overview of the rule you want to create and click **Save** when you are ready to create the rule. When the rule is saved, you see the post success screen. You can choose to **View compliance rules**, **View Compliance**, or **Create another Compliance rule** from this page.

BD to EPG Relationship Configuration

With this feature, you can specify a BD selector to have a fixed number of EPGs. By configuring a BD compliance rule, you can set the maximum number of EPGs with which a BD can be associated.

As a result of this compliance rule, when the requirement set is not satisfied, a violation anomaly will be raised. If the requirement is satisfied, it will raise an enforcement anomaly. Only when the BD selector is not resolved, a warning anomaly will be generated.

The user can configure a requirement to verify that a specified number of EPGs are being associated

with a BD. The supported operators for this requirement are **At least /At most /Equal to**. As an example, if a requirement is configured that the BD must have at least 5 EPGs associated, violation anomalies will be raised if the BD has less than 5 EPGs (0-4). However, if the BD has ≥ 5 anomalies, then an enforcement anomaly will be raised.

Matching Criteria

The following table lists the various objects which are available as matching criteria for a selected object type.

Object Type	Matching Criteria Object
EPG	<ul style="list-style-type: none"> ▪ Tenant ▪ VRF ▪ BD ▪ EPG ▪ App profile ▪ L3 Out ▪ L3 InstP ▪ L2 Out ▪ L2 InstP
Tenant	<ul style="list-style-type: none"> ▪ Tenant
BD	<ul style="list-style-type: none"> ▪ Tenant ▪ VRF ▪ BD
VRF	<ul style="list-style-type: none"> ▪ Tenant ▪ VRF
Contract	<ul style="list-style-type: none"> ▪ Tenant ▪ Contract
Subject	<ul style="list-style-type: none"> ▪ Tenant ▪ Subject
Filter	<ul style="list-style-type: none"> ▪ Tenant ▪ Subject ▪ Filter

The following table lists how the various matching criteria objects can be defined.

Matching Criteria Object Type 1	How to define
Tenant	tn - operator <i>value</i> Object type 2 (Could be either VRF or BD) a. If you select VRF, the rule is further defined as tn - operator <i>value</i> ctx - operator <i>value</i> a. If you select BD, the rule is further defined as tn - operator <i>value</i> bd - operator <i>value</i>
VRF	tn - operator <i>value</i> ctx - operator <i>value</i>
BD	tn - operator <i>value</i> bd - operator <i>value</i>
EPG	tn - operator <i>value</i> ap - operator <i>value</i> epg - operator <i>value</i>
App Profile	tn - operator <i>value</i> ap - operator <i>value</i>
L3 Out	tn - operator <i>value</i> out - operator <i>value</i>
L3 InstP	tn - operator <i>value</i> out - operator <i>value</i> instp - operator <i>value</i>
L2 Out	tn - operator <i>value</i> l2out - operator <i>value</i>
L2 InstP	tn - operator <i>value</i> l2out - operator <i>value</i> instp - operator <i>value</i>
Contract	tn - operator <i>value</i> brc - operator <i>value</i>
Subject	tn - operator <i>value</i> brc - operator <i>value</i> subj - operator <i>value</i>
Filter	tn - operator <i>value</i> fit - operator <i>value</i>



operator and *value* can be set to anything.

You can use the following operators for the custom definitions.

Operator	Description
Must Equal to	This operator returns an exact match of the specified value.
Must Not Equal to	This operator returns all that do not have the same value.
Must Contain	This operator returns all that contain the specified value.
Must not contain	This operator returns all that do not contain the specified value.

Operator	Description
Must begin with	This operator returns all that begin with the specified value.
Must end with	This operator returns all that end with the specified value.
Must not begin with	This operator returns all that do not begin with the specified value.
Must not end with	This operator returns all that do not end with the specified value.

Compliance Rules

Once you create compliance rules, you can generate the Compliance Report to check how much the sites and networks align to the rules.

Click **Manage > Rules > Compliance Rules**. This is where all the created rules are listed.

The Compliance Rules page allows you to view all the rules created in one place. You can perform the following actions on this page:

- Edit or Delete a rule with the "..." button
- Select multiple rules by clicking the checkbox and delete/edit them collectively
- Create a new rule from the **Create Compliance Rule** button.
- Filter the rules using search by the following attributes:
 - Name
 - Description
 - Rule Type
 - State
 - Last Modified Time
- Click on any rule to view the slide-in that brings up the rule summary. It displays the following information:
 - General - Rule description, Site, and State
 - Settings - Rule type, objects used to create the rule, and the configuration compliance rules used.
- **Actions** allows you to edit, delete and disable the rule.

Guidelines and Limitations

If at any point of time after you generate anomalies, one of the following modifications is made:

- The name of an existing compliance rule is edited and modified.
- An existing compliance rule is deleted and a new compliance rule with the same name is created.

- An existing compliance rule is renamed (ABC is renamed to XYZ), and a new compliance rule with the former compliance rule name is created (example, ABC).

To view updated and accurate details in the **Compliance** page, you must regenerate anomalies. In the **Compliance** page, if the **Last Modified** field shows a date and time that is after the snapshot timestamp, anomalies are not generated based on the content in **Compliance** page.

Compliance Analysis

Navigate to **Analyze > Analysis Hub > Compliance**.

- Select a site from the dropdown menu.
- Select the date for which you would like to see the report.

The Compliance Analysis will internally trigger assurance analysis and generate compliance anomalies.

A banner will be displayed if any rule has been modified or a new rule has been added. You can re run the analysis for updated data. A 'modified' or 'new' tag will appear under any rule that has been recently modified or added.



The **Compliance Rules** table and **Anomalies from Violations** table will not be available for reports generated prior to release 6.4(1). You will have to run the analysis again to view the tables.

The **Actions** button allows you to re run the analysis.

The **Summary** displays the number of violations, the top rules by anomaly count, the anomalies from violations and the violations by rule type. You can click on any of the rules in 'Top rules by Violation' to view more details and click the count under 'Number of anomalies from violations' to view the list of anomalies.

The **Anomalies from Violations** lists all the anomalies that were triggered by the rules created. Click any rule in the 'Grouped' view to see the list of anomalies categorized under that group. If you click any rule in the 'Ungrouped' view, you will be redirected to the compliance rule detail page. This can be listed in a group view for all sites or individual view for a specific site. The table lists the severity level of the anomaly, the type of rule that triggered the anomaly, the detection time, and the status.

When you click any rule, it takes you to a slide-in that gives you a summary of your rule (**What's wrong, What triggered this anomaly, What's the impact?, How do I fix it?**).

Use search to filter by the following attributes:

- App Profile Name
- BD Name
- Category
- Compliance Object Name

- Compliance Object Type
- Contract Name
- EPG Name
- Filter Name
- L2 Out Name
- L3 Out Name
- Level
- Rule Name
- Subject Name
- Tenant Name
- VRF Name

The gear icon is used to customize the columns in the table.

The **Compliance Rules** table shows a summary of the rules enforced and violated along with the number for each rule type. The table lists all the rules used to generate the current report. The table specifies whether it's a configuration rule or a communication rule and the number of anomalies from violations for each rule.

Use search to filter by the following attributes:

- Name
- Rule Type
- Enforcement Status
- Verified

The **Create Compliance Rule** button takes you to the rule creation page.



Compliance report is generated once every two hours.

Policy CAM

About Policy CAM

The Policy CAM feature determines how and where resources in the fabric are used. Policy CAM provides information about the resource utilization in the network, and the amount of policy content-addressable memory (Policy CAM) utilization.

Navigate to **Analyze > Analysis Hub > Policy CAM**.

Once you get to Policy CAM, select a site, choose the appropriate snapshot of time within which to view the resource utilization, and click **Apply**.



Within the time range you selected, the last snapshot is considered for each of the site/s included in the site. Therefore, you get the latest state of the application within the selected time range.

Analyze > Analysis Hub > Policy CAM Analyzer

Policy CAM Analyzer

ACI-Paris | Feb 15th 2024, 6:02 AM

Associated Policies
Selecting any combination of objects or policies will update the tables below

- Provider Tenant** View All (10)
- Consumer Tenant** View All (10)
- Provider EPG** View All (72)
- Consumer EPG** View All (44)

openshift462 708 of 1.5 K Entries	openshift462 708 of 1.5 K Entries	aci-containers-nodes 684 of 1.5 K Entries	aci-containers-default 226 of 1.5 K Entries
openshift414 519 of 1.5 K Entries	openshift414 524 of 1.5 K Entries	aci-containers-nodes 320 of 1.5 K Entries	aci-containers-system 198 of 1.5 K Entries
common 69 of 1.5 K Entries	pcv_prod_pci_tn 62 of 1.5 K Entries	aci-containers-system 108 of 1.5 K Entries	aci-containers-system 160 of 1.5 K Entries

- Contract** View All (51)
- Filter** View All (38)
- Node** View All (2)

Policy CAM Statistics

All Policy CAM Rules by Hit Count by EPGs Tenants Leafs Contracts Filters

Filter by attributes

Provider EPG	Consumer EPG	Leaf	Contract	Filter	Consumer VRF	Action
aci-containers-nodes	aci-containers-default	leaf-102 ACI-Paris	openshift-monitoring	openshift-monitoring-filter	openshift-external	✔ Permit
aci-containers-nodes	dom34	leaf-102 ACI-Paris	openshift-monitoring	openshift-monitoring-filter	openshift-external	✔ Permit
aci-containers-nodes	aci-containers-system	leaf-102 ACI-Paris	openshift-monitoring	openshift-monitoring-filter	openshift-external	✔ Permit
aci-containers-nodes	annoateddom	leaf-102 ACI-Paris	openshift-monitoring	openshift-monitoring-filter	openshift-external	✔ Permit
aci-containers-nodes	annoateddom	leaf-101 ACI-Paris	openshift-monitoring	openshift-monitoring-filter	openshift-external	✔ Permit
aci-containers-nodes	dom34	leaf-101 ACI-Paris	openshift-monitoring	openshift-monitoring-filter	openshift-external	✔ Permit

Policy CAM Analyzer displays the following information:

- Associated Policies
- Policy CAM Statistics
- Policy CAM Rules
- All Anomalies



In Nexus Dashboard Insights release 6.3.1.15, Policy CAM is not supported on Cisco Nexus 9000 FX3 switches. In Nexus Dashboard Insights release 6.3.1.40 and later, Policy CAM is supported on Cisco Nexus 9000 FX3 switches.

Associated Policies

Associated policies lists the various objects or policies available. When the policies are viewed in a top to down manner, the lists start with the node that has the maximum utilization followed by the next lower utilization. Each item in each column can be selected to show relevant associations and relationships between the tenants, contracts, and EPGs.

Click **View All** to view all the nodes for the selected object in a side panel.

The following objects or policies are available:

- Provider Tenant
- Consumer Tenant
- Provider EPG
- Consumer EPG
- Contract
- Filter
- Node

Click any of the objects to show all related objects and policies.

Policy CAM Statistics

The policy CAM statistics displays all the nodes and associated rules, and you can drill into details for a specific node here. Click the checkboxes for objects you want to see in the table.

The following objects are available:

- EPGs
- Tenants
- Leafs
- Contracts
- Filters

You can filter the table based on the following attributes:

- Provider EPG
- Consumer EPG
- Leaf
- Contract
- Filter
- Consumer VRF
- Action

The table also shows the hit count in the following timeline:

- 1 month
- 1 week
- 1 hour
- Cumulative

The gear icon allows you to toggle columns to customize the table as per your view.

Policy CAM Rules

In the **Policy CAM Rules** table, you can view the listings for all of the nodes based on the selected snapshot.

You can filter the table based on the following attributes:

- Leaf
- Provider EPG
- Consumer EPG
- Contract
- Filter
- Rule
- Provider Tenant name
- Consumer Tenant name
- Consumer VRF

The following details are available in the Rules table:

- Leaf
- Provider EPG
- Consumer EPG
- Contract
- Filter
- Rule
- Valid Hardware Entry Count

- Provider Tenant name
- Consumer Tenant name
- Consumer VRF

The gear icon allows you to toggle columns to customize the table as per your view.

All Anomalies

In the **Anomalies** table, you can view the anomalies that are generated in the selected snapshot of time, individually by nodes or as an aggregate.

You can filter the anomalies based on the following attributes:

- Anomaly Level
- App Profile Name
- Attachable Access Entity Profiles name
- BD Name
- Concrete Device
- Concrete Interface
- Consumer App Profile Name
- Consumer EPG name
- Contract
- Contract Name
- Device Cluster
- Device Cluster Interface
- Device Selection Policy
- EPG name
- Encap VLAN
- Fabric IP
- Filter name
- Interface Policy Group Name
- Internal/External
- L2 Out Name
- L3 Out Name
- Leaf Interface Profile Name
- Leaf Profile Name
- Logical Interface Context
- Physical Domains Name
- Provider App Profile Name
- Provider EPG Name

- Provider Tenant Name
- Rule Name
- Sites
- Spine Name
- Tenant Name
- Virtual Port Channel

Pre-Change

Pre-Change Analysis

Navigate to **Analyze > Analysis Hub > Pre-Change**.

Pre-Change Analysis allows you to change a configuration for a site, to model the intended changes, perform a Pre-Change Analysis against an existing base snapshot in the site, and verify if the changes generate the desired results.

The screenshot shows the 'Pre-Change Analysis' page in a web application. At the top, there is a breadcrumb trail: 'Analyze > Analysis Hub > Pre-Change Analysis'. A blue button labeled 'Create New Pre-Change Analysis' is in the top right corner. Below the breadcrumb is a search bar labeled 'Filter by attributes'. The main content is a table with the following columns: 'Pre-Change Analysis Name', 'Assurance Entity Name', 'Base Epoch', 'Analysis Status', 'Analysis Submission Time', and 'Submitter ID'. There are two rows of data, both with a 'Completed' status. The first row has 'test2' as the name and 'ACI-Paris' as the assurance entity name. The second row has 'test' as the name and 'ACI-Paris' as the assurance entity name. At the bottom of the table, it says '2 items found' and 'Rows per page 10' with a pagination control showing page 1 of 1.

<input type="checkbox"/>	Pre-Change Analysis Name	Assurance Entity Name	Base Epoch	Analysis Status	Analysis Submission Time	Submitter ID	
<input type="checkbox"/>	test2	ACI-Paris	February 01, 2024, 11:24:35 PM	Completed	February 01, 2024, 11:48:29 PM	Local: admin	...
<input type="checkbox"/>	test	ACI-Paris	February 01, 2024, 11:24:35 PM	Completed	February 01, 2024, 11:43:47 PM	Local: admin	...

© Cisco Systems, Inc.
Current date and time is: Thursday, February 15, 07:50 AM (IST)

After you model the changes for a Pre-Change Analysis job, you can choose **Save** or **Save And Analyze**. By choosing **Save**, you can save the Pre-change Analysis job without having to start the analysis right away. You can return to the job later, edit the changes if required, and then run the analysis later. The **Save** option is supported only for a Pre-Change Analysis job with manual changes.

If you choose **Save And Analyze**, the job gets scheduled and an analysis is provided. The changes are applied to the selected base snapshot, the analysis is performed, and results are generated. For every pre-change analysis job listed in the table, a delta analysis is performed between the base snapshot and the newly generated snapshot.

In Pre-Change Analysis, to see the details of a completed Pre-Change Analysis job, click that job in the table. This opens a new page that displays the following information:

- Dashboard
- Delta Analysis
- Compliance Analysis

test2

Pre-Change Analysis information is based on the simulation created for the Feb 01, 2024, 11:24 PM snapshot.

[Explore Pre-Change Analysis](#)[Actions](#) [Dashboard](#) [Delta Analysis](#) [Compliance Analysis](#)**General Information**

Site
ACI-Paris

Snapshot
02/01/2024 11:24:35 PM

Description (Optional)
Unspecified

Change Definition
Manual

Change Type
ADD

Object Type
fvBD

Bridge Domain's Parent
uni/tn-OOB_Management

Bridge Domain (BD-)
bd1

Private Network

Optimize Wan Bandwidth between sites
no

ARP Flooding
no

Description
-

rogue exception mac wildcard support for bd
no

Clear Endpoints
no

General Information shows the following data:

- Site name
- Snapshot details
- Description
- Change Definition

In case of Manual Changes, you see list of changes that were modeled for that job(Change type, Object type, name alias, Priority, Description, App profile) and in case of JSON/ XML file upload you see the Change Simulation.

As the job is complete, the severity area displays the anomalies that are generated for these changes. To understand the data displayed under Delta Analysis, view [Delta Analysis](#).

To understand the data displayed under Compliance Analysis, view [Compliance](#).

The ... button allows you to perform the following actions:

- Edit Pre Change Analysis
- Clone Pre Change Analysis
- Delete Pre Change Analysis

You can also perform these actions by clicking the checkbox for the desired job or by using the **Actions** button.

Things to remember while performing the three actions:

1. You can clone Pre-Change Analysis jobs for manual changes only.
2. You can delete up to 10 Pre-Change Analysis jobs at a time. You cannot delete a job in the

Running state. If you attempt to do that, an appropriate notification will display.

If anomalies are raised in the analysis, make the required modifications based on the results and re-run the analysis until you obtain satisfactory results. The download option in a Pre-Change Analysis job allows you to download a JSON file that can be uploaded to Cisco APIC. However, if you choose the file upload approach, you can upload a JSON or an XML Cisco APIC configuration file to run a Pre-Change Analysis job.

Once the analysis starts, the status of the job will be shown as Running. During this time, the specified changes will be modeled on top of the base snapshot, and complete logical checks will be run, including Policy Analysis and Compliance. No switch software or TCAM checks will be performed. The status of the Pre-Change Analysis job is marked **Completed** when the entire analysis including Delta Analysis completes. The Delta Analysis is automatically triggered and the associated Pre-Change Analysis job is displayed as running during that time. The Delta Analysis is performed only on checks supported in Pre-Change Analysis job.

You can view changes applied by a user to a specific Pre-Change Analysis job by clicking the job in the table. If the changes are applied manually, you can view the different changes selected by the user. If the job is created using a JSON file, the Change Definition field displays the name of the JSON file from where the changes were imported.

Pre Change Analysis lists all the analyses performed in a tabular form with the following fields:

1. Analysis Name
2. Assurance Entity Name
3. Base Epoch
4. Analysis Status
5. Submitter ID

Pre-Change Analysis Options

The following list specifies the options you can choose on your pre-change analysis job. Only the objects listed are supported.

1. Add, modify, or remove Tenant.
2. Add, modify, remove App EPG (*supported attributes*: preferred group member, intra EPG isolation; *relations for App EPG*: BD, provided, consumed and taboo contracts; *export/import of contracts* is not supported.)
3. Add, modify, or remove a VRF (*supported attributes*: policy control enforcement preference, policy control enforcement direction, BD enforcement status, preferred group member, description).
4. Add, modify, or remove a BD (*supported attributes*: description, optimize WAN bandwidth, type, ARP flooding, IP learning, limit IP learning to subnet, L2 unknown unicast, unicast routing, multi-destination flooding, multicast allow, and L3 unknown multicast flooding).
5. Add, modify, or remove a contract (*supported attributes*: scope, description).
6. Add, modify, or remove a contract subject (*supported attributes*: reverse filter ports, description, priority, target DSCP, filter name, forward filter name, and reverse filter name).

7. Add, modify, or remove subnets (*supported attributes*: scope, preferred, description, primary IP address, virtual IP address, and subnet control).
8. Add, modify, or remove an App profile (priority, description).
9. Add, modify, or remove an L3Out (*supported attributes*: description, VRF name, Target DSCP, and route control enforcement).
10. Add, modify, or remove an L2Out (*supported attributes*: description, BD name, encapsulation type, and encapsulation ID).
11. Add, modify, or remove an L3 Ext EPG (*supported attributes*: preferred group member, description, priority; *supported relations*: VRF, provided contracts, consumed contracts, taboo, and target DSCP).
12. Add, modify, or remove an L2 Ext EPG (*supported attributes*: preferred group member, description, priority, target DSCP and provided contracts, supported contracts, and taboo contracts).
13. Add, modify, or remove L3 Ext EPG Subnets (*supported attributes*: description, and scope).
14. Add, modify, or remove a Taboo Contract (*supported attributes*: description).
15. Add, modify, or remove a Taboo Subject (*supported attributes*: name, description; *supported relations*: vzRsDenyRule).
16. Add, modify, or remove a Filter and Filter entries.

For Fabric Access Policies, you can choose to add the following to your pre-change analysis job:

1. Add, modify, or remove relationship between EPG and a physical domain.
2. Add, modify, or remove relationship between physical domain and a corresponding VLAN pool.
3. Add, modify, or remove relationship between physical domain and Attachable Entity Profile.
4. Add, modify, or remove a leaf interface profile.
5. Add, modify, or remove a port selector.
6. Add, modify, or remove a switch profile.
7. Add, modify, or remove a switch selector.
8. Add, modify, or remove an interface policy group.
9. Add, modify, or remove an interface policy for CDP and LLDP.

Guidelines and Limitations

When using Pre-Change Analysis follow these guidelines and limitations:

- Pre-change Analysis can be conducted for sites and uploaded files.
- More than one Pre-change Analysis can be run on the same base snapshot.
- Pre-Change Analysis cannot be run for a pre-change snapshot being used as a base snapshot.
- Only logical configuration anomalies are modeled and run in a Pre-Change Analysis. Switch software and TCAM changes are not modeled. After the analysis completes, a Delta Analysis will automatically start to compare the snapshot, generated due to the Pre-Change Analysis, with the base snapshot. Delta Analysis is performed only on checks supported in the Pre-Change Analysis job.

- During a pre-change analysis, certain anomalies that exist in the base snapshot will not be analyzed in the pre-change analysis. As a result, these anomalies will not appear in the Pre-Change Analysis snapshot even though the violation continues to exist. The reason that such an event is not analyzed in a pre-change analysis is because these anomalies require not just logical data, but they also require switch software and TCAM data.
- Compliance Analysis displays the results of compliance checks in the Pre-Change Analysis snapshot.
- A local search of anomalies from a Pre-Change Analysis snapshot can be performed and viewed in the results section by navigating to specific tabs for **Dashboard**, **Delta Analysis**, **Compliance Analysis**, and **Explore**.
- Pre-Change Analysis does not support or analyze any service chain related changes or objects.
- Delta Analysis does not allow a Pre-Change Analysis snapshot to be selected.
- If configuration data does not exist for a base snapshot, and you run a pre-change analysis job using this snapshot, new logical configuration files will not be generated. For such pre-change analysis jobs, the Download icon will be grayed out/disabled in the side panel. You will not be able to download a new logical configuration.
- The Pre-Change Analysis could go into a Failed state if an imported configuration has unsupported objects. Figure out the Cisco ACI objects that are unsupported by referring to the [Pre-Change Analysis Options](#) section, remove them, and import the configuration again before starting another Pre-Change Analysis job. If there is a failed Pre-Change Analysis, the error message for the failure is displayed in the Pre-Change Analysis table under **Analysis Status**.
- The Pre-change Analysis feature is supported in Cisco APIC release 3.2 or later. If you attempt to run a Pre-change Analysis with a Cisco APIC release earlier than release 3.2, an ERROR message indicates that Pre-Change verification is supported on APIC 3.2 or higher, and you cannot run the analysis.
- If there is an analysis that is currently running when you start a Pre-Change Analysis, that job is completed first. The new jobs are serviced in the order the jobs are scheduled. Cisco Nexus Dashboard Insights runs the jobs in the order that best suites the schedule and the available resources. All jobs, including the Pre-Change Analysis job are given the same priority.
- Currently, you can upload a JSON or an XML Cisco APIC configuration file to run a Pre-Change Analysis job. The maximum **change file size** supported are: 10 MB for vND and 50 MB for pND. An uploaded file will be pruned by removing white spaces and endpoint objects (fvCEp) to reduce the file size.
- You can save as many Pre-Change Analysis jobs as you want. However, for a site, you can only run a Pre-Change Analysis job one at a time.

Support for Multiple Objects in Pre-Change Analysis

In addition to multiple tenants, you can also add multiple infrastructure objects as part of a Pre-Change Analysis JSON or XML job. The Pre-Change Analysis upload path allows you to add, modify, and delete multiple objects across the policy universe. There are no additional configurations required to use this feature. Your Pre-Change Analysis job for multiple objects will run, based upon the file/s you upload.

The following file upload formats are accepted:

- A JSON or XML file with IMDATA of size 1.

- An IMDATA that contains a single subtree of the intended changes. The root of the subtree can be the UNI or any other Managed Object as long as the changes are represented as a single subtree.
- Use the file that you had uploaded from a JSON or XML path to perform a Pre-change Analysis. After the Pre-Change Analysis is complete, you can upload the same file to ACI to be used to make the changes.

Known Issues for Pre-Change Analysis

- When Pre-Change Analysis scale limits are exceeded, the analysis can fail with no error message.
- For Pre-Change Analysis jobs, you must not modify configurations where the total number of EPGs, BDs, VRFs are greater than 16,000.
- When creating a new Pre-Change Analysis, note the following:
 - If the JSON/XML file size being uploaded is less than 100 MB but greater than 15 MB, then the API validates the file and throws a validation error as follows: *Uploaded file size exceeds the 15MB(pND)/8MB(vND) maximum limit.* When users access Cisco Nexus Dashboard Insights, and try to create a Pre-Change Analysis job with a file size greater than 15MB(pND)/8MB(vND), the UI throws the following error: *File size cannot be larger than 15MB(pND)/8MB(vND).* Therefore, files larger than 15MB(pND)/8MB(vND) are not supported in Pre-Change Analysis.
 - If you upload a file with unsupported objects, Cisco Nexus Dashboard Insights will remove the unsupported object and run the job.
- A Pre-change Analysis job may fail or return incorrect results if the Cisco ACI configuration has features that are unsupported by Cisco Nexus Dashboard Insights.
- Pre-change Analysis is not supported in Cisco ACI configurations that contain service chains.
- Cisco Nexus Dashboard Insights performs a limited set of checks on the JSON file uploaded for pre-change analysis. Cisco ACI may reject this file.
- Pre-change Analysis may incorrectly report errors for attributes of subnets of external routed networks.
- Pre-change Analysis is supported in the following Cisco APIC releases:
 - For 3.2(x) release, 3.2(9h) and earlier are supported
 - For 4.0(x) release, 4.0(1h) and earlier are supported
 - For 4.1(x) release, 4.1(2x) and earlier are supported
 - For 4.2(x) release, 4.2(7s) and earlier are supported
 - For 5.0(x) release, 5.0(2e) and earlier are supported
 - For 5.1(x) release, 5.1(4c) and earlier are supported
 - For 5.2(x) release, 5.2(4d) and earlier are supported
 - For 5.3(x) release, 5.3(1b) and earlier are supported
 - For 6.0(x) release, 6.0(4c) and earlier are supported

Create Pre-Change Analysis Job

1. Navigate to **Analyze > Analysis Hub > Pre-Change**.

2. In **Pre-Change**, click **Create Pre-Change Analysis**. In **Create Pre-Change Analysis**, perform the following actions:

General

- a. In the **Pre-Change Analysis Name** field, enter a name.
- b. In the **Description** field, add a description for the analysis if you would like to.
- c. In the **Site** field, choose the appropriate site.
- d. In the **Snapshot** field, specify the appropriate snapshot.

Change

- a. Under **Change**, choose the appropriate option. (**Import JSON/XML File** or **Manual Changes**).



Depending upon your selection, the relevant fields are displayed for you to populate.

If you choose the file import option to upload a JSON or XML file upload, you must click **Save & Run** to start the Pre-Change Analysis operation.

If you choose the manual changes option, select the **Change Type** and the **Object Type** and then you can either save & run the job, or save the job to start it at a later time by clicking **Actions > Edit Pre-Change Analysis** and clicking **Save & Run**. When in **Edit**, you can also change some of the fields if required.

Complete the selections as appropriate, and click **Save** or **Save & Run**.

After a Pre-Change Analysis job is completed, the **Pre-Change Analysis** table displays the status for the job as completed.

Click the Pre-Change Analysis Name for which you want to view the details. In a sidebar to the right, the details are displayed in a column including the general information such as the name of the job, snapshot, and change definition type. The list of changes modeled for the job are also available. If you are viewing a completed job, the anomalies that were generated as a result of the changes are displayed at the top of this page.

For completed jobs, click the icon on the top right of the sidebar to navigate to the results page. Further details about the job are available here under the specific tabs for **Dashboard**, **Delta Analysis**, **Compliance Analysis**.

Download Pre-Change Analysis Job

You can download an existing Pre-Change Analysis as follows:

- In the **Pre-Change Analysis** table, click the appropriate pre-change analysis name for a completed Pre-Change Analysis job. Click the download icon to download the file.
- The pre-change analysis downloads as an offline tar file with the pre-change analysis contents displayed in JSON format.



In the downloaded file, you can view all the attributes which include attributes that

are modified and those that are not modified. If desired, the downloaded file can be uploaded to your Cisco APIC.

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: <http://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)

© 2017–2024 Cisco Systems, Inc. All rights reserved.