



Analyzing and Troubleshooting Your Network, Release 4.3.1

Table of Contents

New and Changed Information	1
Analyzing and troubleshooting your network using Analysis Hub	2
Gathering fabric resource information using Policy CAM	3
Associated policies	5
Policy CAM Statistics	6
Policy CAM Rules	7
All Anomalies	8
Microsegmentation endpoint groups (uSeg EPG) and PBR assurance	9
Defining communication and configuration rules using Compliance (ACI fabrics)	11
Types of compliance	11
Communication	11
Configuration	11
Examples of compliance	12
Integrating Splunk for advanced monitoring and compliance	18
Compliance rules	19
Interpretation of compliance rules	20
Compliance analysis	20
Compliance anomalies	21
Guidelines and limitations for compliance	21
Guidelines for compliance	21
Guidelines for creating a compliance rule	22
Guidelines for import configuration compliance	22
Guidelines for compliance rule naming	22
Verified scalability guidelines for template-based compliance	22
Guidelines for template-based compliance	22
Limitations for compliance analysis	23
Limitations for Integrating Splunk for advanced monitoring and compliance	23
Create compliance communication rule	23
Create compliance rule with snapshot selection	24
Import configuration compliance	24
Create compliance rule with manual configuration	24
Create template-based compliance	25
Trigger a compliance analysis	25
Templates to configure object selectors	25
Selection, status, and match criteria	25
Tags and annotations	30
Communication compliance	32
Communication types	32
Communication type and traffic selector rules selections with the resultant compliance rule type	33
Traffic selector rules available	33

Manual configuration compliance	34
Attribute requirement that can be set according to the objects selected	34
BD to EPG relationship configuration	35
Matching criteria	35
Objects available as matching criteria for a selected object type	35
Define matching criteria objects	36
Operators for custom definitions	37
Understanding the hardware and software lifecycles for your network using Conformance	38
Access conformance report	39
View conformance report	39
Analyzing endpoint flows using Connectivity Analysis	41
Connectivity Analysis support for Layer 4 to Layer 7 services	41
Explorer integration with connectivity analysis	42
Silicon One Packet Tracer integration	42
Prerequisites for L4-L7 Connectivity Analysis	42
Guidelines and limitations for Connectivity Analysis	43
Connectivity Analysis options	43
Layer 2 ToR support	44
Guidelines and limitations (connectivity analysis): For all fabrics	45
Guidelines and limitations (connectivity analysis): For fabric groups	45
Guidelines and limitations (connectivity analysis): For Cisco ACI fabrics	45
Guidelines and limitations (connectivity analysis): For standalone NX-OS fabrics	46
Supported topologies	46
Create connectivity analysis	47
View connectivity analysis	49
Manage connectivity analysis	56
Filtering information	57
Monitoring the latency, congestion, and drops for your network, using traffic analytics	59
Traffic analytics conversations	61
Traffic analytics transit conversations	61
Upgrade behavior	61
Conversation tagging and statistics	62
Enhancements to traffic analytics endpoints view	63
Traffic analytics for multi-fabric and external conversations	63
Traffic analytics support for traffic redirection to L4-L7 services	64
Traffic analytics support for L3 external filter rules	64
Guidelines and limitations for traffic analytics	65
General guidelines for traffic analytics	65
Support information for traffic analytics with any fabric type	65
Guidelines and limitations for traffic analytics compatibility mode	66
Support information for traffic analytics with Cisco ACI fabrics	67
Limitations for traffic analytics with Cisco ACI fabrics	67
Support information for traffic analytics with NX-OS fabrics	67

Limitations for traffic analytics with NX-OS fabrics	68
Limitations for traffic analytics with multi-fabric and external conversations	68
Limitations for traffic analytics with fabric groups	68
Traffic analytics support for Silicon One (S1) platforms	69
Enhancements to Traffic analytics compatibility mode	69
Supported switches and behavior	70
Supported topologies	72
Upgrade and migration	72
Guidelines and limitations for Traffic analytics support for S1 platforms	72
Configure traffic analytics flow collection	73
Apply the traffic analytics configuration	73
Apply the traffic analytics compatibility mode configuration	74
View traffic analytics	75
View traffic analytics for an individual fabric	75
View traffic analytics for a fabric group	78
View traffic analytics status for all fabrics	80
Manage service categories	80
Apply filters to restrict traffic analytics for traffic analytics L3 external filtering	82
Create an interface filter rule	83
View traffic analytics for endpoints	85
View traffic analytics for service endpoints	86
Find traffic analytics conversations for an endpoint in a fabric	87
Find traffic analytics conversations for an endpoint in a fabric group	90
Use the flow troubleshoot workflow	91
Exploring the energy usage for your network using Energy Management	95
Cisco Energy Manager	96
View the energy management report for switches	96
View the energy management report for PDUs	99
Comparing the configurations between two time periods for your fabric using Delta analysis	102
Health Delta	102
Policy Delta for ACI	102
Policy Delta for standalone NX-OS	102
Guidelines and limitations for delta analysis	103
Create delta analysis	103
View delta analysis	104
View health delta analysis	104
View policy delta analysis	106
Viewing the impact of configuration changes using pre-change analysis (ACI fabrics)	109
Pre-change analysis options	111
Guidelines and limitations for pre-change analysis	112
Support for multiple objects in pre-change analysis	113
Known issues for pre-change analysis	114
Create pre-change analysis job	114

Download pre-change analysis job	115
Collecting and analyzing your device logs using Log Collector	116
Device connectivity notifier for TAC-initiated collector	116
Guidelines and Limitations	116
Log collector dashboard	116
TAC-initiated log collector	117
Upload logs to Cisco Intersight Cloud	118
Guidelines and Limitations	118
Collecting information on bugs that might affect your network using Bug scan	119
Enhancements to Bug scan	119
Bug scan schedule	120
Guidelines and limitations	120
Enable Bug scan	121
Enable Bug scan using Cisco Intersight account	121
Upload metadata using air-gapped mechanism	122
View known bugs	122
View known bugs for an individual fabric	124
Tracking endpoints within a data center using Endpoint Locator (NX-OS fabrics)	126
Backup and Restore	128
EPL Connectivity Options	128
Nexus Dashboard Cluster Mode: Physical Server to VM Mapping	128
Configuring Endpoint Locator	128
Configure Endpoint Locator from Nexus Dashboard user interface	129
Configuring Endpoint Locator for Single VXLAN EVPN Site	130
Configuring Endpoint Locator for interconnected VXLAN fabrics	133
Configuring Endpoint Locator for vPC Fabric Peering Switches	136
Configuring Endpoint Locator for External Fabrics	137
Configuring Endpoint Locator for eBGP EVPN Fabrics	137
Flushing the Endpoint Database	138
Monitoring Endpoint Locator	138
Disabling Endpoint Locator	139
Copyright	140

New and Changed Information

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

Release Version	Feature	Description
Nexus Dashboard 4.3.1	Connectivity Analysis support for fabric groups	Connectivity analysis is now available at the fabric group level. For more information, see Analyzing endpoint flows using Connectivity Analysis .
Nexus Dashboard 4.3.1	Fabric agent package updates in monitor mode	Connectivity Analysis can update fabric agent packages for standalone NX-OS fabrics monitor mode when the required telemetry, license, metadata, and image management prerequisites are met. For more information, see Create connectivity analysis .
Nexus Dashboard 4.3.1	Traffic Analytics support for fabric groups	Traffic Analytics is now available at the fabric group level. For more information, see View traffic analytics for a fabric group .
Nexus Dashboard 4.3.1	Traffic Analytics support for Silicon One (S1) platforms	Beginning with Nexus Dashboard 4.3.1, Nexus Dashboard now supports Traffic analytics compatibility in Netflow mode for Silicon One (S1) switches running NX-OS 10.6(3) release or later. This release also supports traffic analytics for sFlow-based N9K-C9804, N9K-C9808, and G200 switches. • NX-OS fabrics support traffic analytics compatibility in both Netflow and sFlow modes, depending on the supported switch models. • ACI fabrics will only support Traffic analytics compatibility in Netflow mode. • Traditional fabrics with only Cloud Scale (CS) switches will continue to support Traffic analytics compatibility in Netflow mode. For more information, see Traffic analytics support for Silicon One (S1) platforms.
Nexus Dashboard 4.3.1	Cisco Nexus 9000 Series name change	Cisco Nexus 9000 Series is now the Cisco N9000 Series.

Analyzing and troubleshooting your network using Analysis Hub

In the Analysis Hub area, you can analyze and troubleshoot your network using advanced analytics tools, which allows you to gain valuable insights into the performance and health of your network.

These sections provide more information on different ways that you can analyze and troubleshoot your network:

- [Gathering fabric resource information using Policy CAM](#)
- [Defining communication and configuration rules using Compliance \(ACI fabrics\)](#)
- [Understanding the hardware and software lifecycles for your network using Conformance](#)
- [Analyzing endpoint flows using Connectivity Analysis](#)
- [Monitoring the latency, congestion, and drops for your network using Traffic Analytics](#)
- [Exploring the energy usage for your network using Energy Management](#)
- [Comparing the configurations between two time periods for your fabric using Delta analysis](#)
- [Viewing the impact of configuration changes using pre-change analysis \(ACI fabrics\)](#)
- [Collecting and analyzing your device logs using Log Collector](#)
- [Collecting information on bugs that might affect your network using Bug scan](#)
- [Tracking endpoints within a data center using Endpoint Locator \(NX-OS fabrics\)](#)

Gathering fabric resource information using Policy CAM

The Policy CAM feature determines how and where Tcam resources in the fabric are used. Policy CAM provides information about the Tcam resource utilization in the network, and the amount of policy content-addressable memory (Policy CAM) utilization. The Policy CAM page displays policies for microsegmentation endpoint groups (uSeg EPGs). You can verify TCAM rules and policy associations.

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

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

With Nexus Dashboard 4.3.1, the Policy CAM feature is supported on VXLAN iBGP and VXLAN external fabrics in addition to ACI fabrics which were already supported in previous Nexus Dashboard releases. In the Nexus Dashboard UI, you can now view associations between Security groups, Contracts, and Filters on NX-OS switches. You can also view the rules created by each Contract and packet hit count for each rule on your fabric.

The Policy CAM feature for NX-OS fabrics requires the VXLAN Group Policy Option (GPO) security group configuration on your fabrics. For more information, see the "Working with security groups" section in [Working with Segmentation and Security for Your Nexus Dashboard VXLAN Fabric](#). For VXLAN external fabrics, you must manually configure the security group configuration to see the Policy associations utilization in the Nexus Dashboard **Analysis Hub** UI.

The Policy CAM feature is supported on these platforms:

- Cisco N9300-FX3, Cisco N9300-GX, and Cisco N9300-GX2 switches with NX-OS Release 10.4(3) and later releases.
- Cisco N9000 FX/FX/H switches with NX-OS Release 10.5.2 and later releases.
- Cisco N9408FX with NX-OS Release 10.5.1 and later releases.

← Analysis hub

Policy CAM analyzer

DC-NX1 | Active now | Re-sync successful | Last Re-sync completed at 11:41 PM on July 28, 2025

Showing report as of Jul 28, 2025, 11:41 PM

Refresh | Re-sync

Associated policies

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

Source security group	Destination security group	Contract	Filter	Switch
DG_vlan96_SC1_S1_1	DG_vlan96_SC1_S1_1	Contract-S1-1000_TCP_port_401	TCP_port_401	blr6-leaf3
DG_vlan96_SC1_S1_10	DG_vlan96_SC1_S1_10	Contract-S1-1001_TCP_port_401	TCP_port_402	blr6-leaf4
DG_vlan96_SC1_S1_100	DG_vlan96_SC1_S1_100	Contract-S1-1002_TCP_port_401	TCP_port_403	

Policy CAM statistics

All Policy CAM Rules by Hit Count by ☒ Security groups ☐ Switches ☐ Contracts ☐ Filters

Filter by attributes

Source security group	Destination security group	Source VRF	Action	1 Month	1 Week	1 Day	2 Hours	Cumulative
SG_vlan95_SC1_S1_10	DG_vlan96_SC1_S1_10	myvrf_80001	Permit	-	-	1025	1025	1025
SG_vlan95_SC1_S1_100	DG_vlan96_SC1_S1_100	myvrf_80001	Permit	-	-	1025	1025	1025
SG_vlan95_SC1_S1_101	DG_vlan96_SC1_S1_101	myvrf_80001	Permit	-	-	1025	1025	1025
SG_vlan95_SC1_S1_102	DG_vlan96_SC1_S1_102	myvrf_80001	Permit	-	-	1025	1025	1025
SG_vlan95_SC1_S1_103	DG_vlan96_SC1_S1_103	myvrf_80001	Permit	-	-	1025	1025	1025
SG_vlan95_SC1_S1_104	DG_vlan96_SC1_S1_104	myvrf_80001	Permit	-	-	1025	1025	1025
SG_vlan95_SC1_S1_105	DG_vlan96_SC1_S1_105	myvrf_80001	Permit	-	-	1025	1025	1025
SG_vlan95_SC1_S1_106	DG_vlan96_SC1_S1_106	myvrf_80001	Permit	-	-	1025	1025	1025
SG_vlan95_SC1_S1_107	DG_vlan96_SC1_S1_107	myvrf_80001	Permit	-	-	1025	1025	1025
SG_vlan95_SC1_S1_108	DG_vlan96_SC1_S1_108	myvrf_80001	Permit	-	-	1025	1025	1025

There are some minor variations in how the latest snapshot report is generated for Cisco ACI and NX-OS fabric types:



- In Cisco ACI fabrics, within the time range you choose, the last snapshot is considered for each of the fabrics included in the fabric. Therefore, you get the latest state of the application within the chosen time range.
- In NX-OS fabrics, when you choose the time range, the report is based on the end time.
- In NX-OS fabrics, the **Resync** button helps to detect any discrepancies between the NX-OS fabric and the intended configuration. This helps to make sure the latest policy changes are detected and reflected in the Nexus Dashboard UI.
- TCAM utilization count for each policy is not supported for NX-OS fabrics in this release.

The **Policy CAM Analyzer** page for Cisco ACI fabrics displays the **Associated Policies** and **Policy CAM Statistics**.

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 Entries	openshift462 Entries	aci-containers-nodes Entries	aci-containers-default Entries
openshift414 Entries	openshift414 Entries	aci-containers-nodes Entries	aci-containers-system Entries
common Entries	pcv_prod_pci_tn Entries	aci-containers-system Entries	aci-containers-system 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



The **All Anomalies** information is not supported for NX-OS fabrics.

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. For Cisco ACI fabrics, each item in each column can be chosen to show relevant associations and relationships between the tenants, contracts, and EPGs. Similarly, for NX-OS fabrics, these columns show the associations and relationships between the Security group, Contracts, and Filters.

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

These objects or policies are available for fabrics.

- Cisco ACI fabrics:
 - Provider Tenant
 - Consumer Tenant
 - Provider EPG
 - Consumer EPG
 - Contract
 - Filter
 - Node
- NX-OS fabrics:
 - Source security group

- Destination security group
- Contract
- Filter
- Switch

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.

These objects are available.

- For Cisco ACI fabrics:

- EPGs
- Tenants
- Leafs
- Contracts
- Filters

- For NX-OS fabrics:

- Security groups
- Switches
- Contracts
- Filters

You can filter the table based on these attributes.

- For Cisco ACI fabrics:

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

- For NX-OS fabrics:

- Source security group
- Source security group policy ID
- Destination security group
- Destination security group policy ID

- Source VRF
- Contract
- Filter
- Action

The table also shows the hit count in these timelines. * For Cisco ACI fabrics: **1 month** 1 week **1 day** Cumulative

- For NX-OS fabrics:

- 1 month
- 1 week
- 1 day
- 2 hours
- 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 chosen snapshot.

You can filter the table based on these attributes.

- For Cisco ACI fabrics:

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

- For NX-OS fabrics:

- Source security group
- Source security group policy ID
- Destination security group
- Destination security group policy ID
- Source VRF
- Switch

- Contract
- Filter
- Rule

These details are available in the Rules table.

- For Cisco ACI fabrics:
 - Leaf
 - Provider EPG
 - Consumer EPG
 - Contract
 - Filter
 - Rule
 - Valid Hardware Entry Count
 - Provider Tenant name
 - Consumer Tenant name
 - Consumer VRF

For NX-OS fabrics:

- Source security group
- Source security group policy ID
- Destination security group
- Destination security group policy ID
- Source VRF
- Switch
- Contract
- Filter
- Rule

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 chosen snapshot of time, individually by nodes or as an aggregate.



The **Anomalies** table is not supported on the NX-OS fabrics.

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
- Fabrics
- Spine Name
- Tenant Name
- Virtual Port Channel

Microsegmentation endpoint groups (uSeg EPG) and PBR assurance

Previously, if you used Cisco Application Centric Infrastructure (ACI)'s Policy-Based Redirect (PBR) functionality with microsegmentation endpoint groups (uSeg EPG) based on IP or MAC addresses previously lacked comprehensive assurance and monitoring capabilities. This feature addresses the

critical need by introducing enhanced assurance, monitoring, and visibility for (uSeg) EPGs and PBR configurations, enabling users to proactively identify and resolve potential issues in their ACI fabrics and supporting the transition to more app-centric deployments with granular security policies.

This feature provides the following enhancements:

- Gain deeper insights into (uSeg) EPG deployments through several enhancements:
 - **Explorer** policy table: A dedicated uSeg tag is added to identify uSeg EPGs.
 - **Flows and Traffic Analytics:**
 - (uSeg) EPG information is displayed in flow and Traffic Analytics conversations.
 - Flow records will carry source and destination class (sclass and dclass) information, which upon enrichment, will provide either ESG or uSeg details.
 - For Traffic Analytics compatibility mode, uSeg information will be enriched from Endpoint Messages, and collector IP databases will contain uSeg information.
 - The existing EPG field in the Flow/Traffic Analytics API will now display either EPG or uSeg.
 - **Endpoint Information** page: (uSeg) EPG details are displayed, along with links to (uSeg) EPG information within the Segmentation and Security section. This page also presents a Traffic Analytics score for the endpoint, providing an immediate health indicator.
- Anomaly checks proactively identify misconfigurations in (uSeg) EPGs, including missing classification, missing static leaf configurations for IP or MAC-based (uSeg) EPGs, and incorrect classification programming on static leaf.
- The existing anomaly for PBR destination reachability is enhanced to verify endpoint learning for PBR destinations when used with (uSeg) EPGs. Existing compliance flows now seamlessly support (uSeg) EPGs.
- The **Policy CAM** page displays policies related to (uSeg) EPGs.
- Delta analysis includes (uSeg) EPGs as part of the EPGs resource in the **Health Delta** page by **Resources** view.
- Connectivity analysis supports (uSeg) endpoints.



Anomaly checks support IP-based or MAC-based (uSeg) EPGs for physical domain associations only and do not support Virtual Machine Manager (VMM) domain or virtual machine (VM) attribute-based classifications.

Defining communication and configuration rules using Compliance (ACI fabrics)

Compliance allows you to define a set of rules to enforce your communication and configuration standards or expectations. Compliance flows and communication and configuration rules support microsegmentation endpoint groups (uSeg EPGs). You can enforce isolation or communication standards for endpoints grouped by IP or MAC addresses.

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

You can change the state of any rule after the rules have been created. If the rule is in the 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.

Go to **Configure > Rules** and enable or disable the specific rule from the Rule State column in the table. 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**.

Types of compliance

There are two compliance types:

- [Communication](#)
- [Configuration](#)

Communication

The Communication compliance type enables communication or isolation between network objects that meet business and regulatory purposes. To create a compliance communication rule, see [Create compliance communication rule](#).

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

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

Configuration Compliance helps perform a configuration compliance check against a specified configuration.

Configuration compliance can be further classified into four types:

- **Snapshot Settings Compliance:** 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. To create a compliance rule with snapshot selection, see [Create compliance rule with snapshot selection](#).
- **Manual Configuration:** You can configure this for certain objects such as BD, VRF, EPG, Contract, Subject, and Filter. All objects types are not supported. To create a compliance rule with manual configuration, see [Create compliance rule with manual configuration](#).
- **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. To create a template-based compliance, see [Create template-based compliance](#).

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.

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.

- **Import Configuration Compliance:** You can perform an import configuration against a specified configuration. You specify a configuration file or snapshot, and Nexus Dashboard 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. To create an import configuration compliance, see [Import configuration compliance](#).

Examples of compliance

- Example of template-based 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**. Select contract subjects which have **name** as any (wildcard) and **nameAlias** must be ABC. The **status MUST_EXIST** means for all the parent nodes that exists, at least one of obj_type must exist. If a select is defined it should obey that condition.

```
{
  "vzBrCP":
  {
    "attributes":
    {
      "STATUS": "MUST_EXIST",
      "SELECT(name)": "REGEX(Ctrct_[1-3])",
    }
  }
}
```

```

    "MATCH(scope)": "EXACT(context)"
  },
  "children":
  [
    {
      "vzSubj":
      {
        "attributes":
        {
          "SELECT(name)": "REGEX(.*)",
          "nameAlias": "ABC"
        },
        "children":
        [
          {
            "vzRsSubjFiltAtt":
            {
              "attributes":
              {
                "SELECT(tnVzFilterName)": "ENDS_WITH(3_1_1)",
                "MATCH(action)": "deny"
              }
            }
          }
        ]
      }
    }
  ]
}

```

- o Does BD has a IPv4 subnet and an L3Out associated with it for specific tenants? This can be evaluated in the below template.

```

{
  "fvTenant":
  {
    "attributes":
    {
      "SELECT(dn)": "OR(uni/tn-mgmt,uni/tn-tcam_comp_aepg_aepg,uni/tn-
Corp102)"
    },
    "children":
    [
      {

```

```

"fvBD":
{
  "attributes":
  {
    "SELECT(name)": "REGEX(.*)"
  },
  "children":
  [
    {
      "fvSubnet":
      {
        "attributes":
        {
          "SELECT(ip)": "REGEX^[A-Fa-f0-9]{1,4}\\\:*)",
          "ctrl": "nd",
          "ipDPLearning": "enabled",
          "scope": "public"
        }
      }
    },
    {
      "fvRsBDToOut":
      {
        "attributes":
        {
          "STATUS": "MUST_EXIST",
          "MATCH(tnL3extOutName)": "REGEX(L3Out_W02_[A-Za-z0-9])"
        }
      }
    },
    {
      "fvRsCtx":
      {
        "attributes":
        {
          "MATCH(tnFvCtxName)": "REGEX(VRF_W02_[A-Za-z]*)",
          "STATUS": "MUST_NOT_EXIST"
        }
      }
    },
    {
      "fvRsBDToNdP":
      {
        "attributes":
        {

```

```

        "STATUS": "MUST_EXIST"
      }
    }
  }
]
}
}
}

```

- o EPG must not has a VMM Domain Configured. This can be evaluated in the below template.

```

{
  "fvTenant": {
    {
      "attributes": {
        {
          "SELECT(dn)": "STARTS_WITH(uni/tn-NAE_contract)"
        },
      },
      "children": [
        {
          {
            "fvAp": {
              {
                "attributes": {
                  {
                    "SELECT(name)": "REGEX(.*)"
                  },
                },
                "children": [
                  {
                    {
                      "fvAEPg": {
                        {
                          "attributes": {
                            {
                              "floodOnEncap": "disabled",
                              "hasMcastSource": "no",
                              "MATCH(name)": "REGEX(^EPG_W02_[A-Za-z0-9_-]*)",
                              "pcEnfPref": "unenforced",
                              "prefGrMemb": "include",
                              "MATCH(prio)": "REGEX(^level[0-9])",
                              "shutdown": "no"
                            },
                          },
                        },
                      },
                    },
                  ],
                },
              },
            },
          ],
        },
      ],
    },
  },
}

```



```

"children":
[
  {
    "fvRsDomAtt":
    {
      "attributes":
      {
        "instrlmedcy": "lazy",
        "reslmedcy": "pre-provision",
        "STATUS": "MUST_NOT_EXIST"
      },
      "children":
      [
        {
          "fvAEPgLagPolAtt":
          {
            "attributes":
            {
              "annotation": " "
            },
            "children":
            [
              {
                "fvRsVmmVSwitchEnhancedLagPol":
                {
                  "attributes":
                  {
                    "MATCH(tDn)":
                    " ENDS_WITH(LACP_SDN)"
                  }
                }
              }
            ]
          }
        }
      ]
    }
  }
]

```

```
}  
}
```

- Invalid Example of template-based compliance
 - 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**. This is a violation because you cannot allow two SELECT criteria to coexist for the same object time because as can lead to two different ways of selection and validation of objects. So **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*)"  
          }  
        }  
      }  
    ]  
  }  
}
```

Integrating Splunk for advanced monitoring and compliance

Splunk is now a native application within the Nexus Dashboard Analysis Hub, providing you with robust capabilities to build personalized dashboards, generate insightful reports, and configure critical alerts. In Nexus Dashboard 4.2.1, native Splunk support is only available on single ND-NODE-G5S node clusters.

This integration enables NetOps, network administrators, and SOC users to proactively monitor, analyze, and remediate network events using AI/ML-driven insights and automation—all from a single platform.

Splunk integration provides these advanced capabilities within the Analysis Hub:

- **Custom dashboards:** Create and view customer-specific dashboards and compliance visualizations.
- **Predefined dashboards:** Predefined dashboards for anomalies, advisories, and audit logs.

Deployment and access

The details below explain how Splunk is deployed and accessed through Nexus Dashboard:

- **Packaging:** Splunk is packaged as a Nexus Dashboard application and is only available in a single node physical deployment for ND-NODE-G5S hardware.
- **Enablement:**
 - Enable Splunk in Nexus Dashboard's advanced settings under **Admin > System Settings > General > Advanced Settings > Edit > Enable Splunk**.
 - A separate Splunk license is required. You can either copy and paste the content of the license file or upload the license file.

The screenshot shows the 'Advanced settings' page in the Nexus Dashboard. The left sidebar contains navigation links: Home, Manage, Analyze, and Admin (selected). The main content area is titled 'Advanced settings' and includes a 'Back' link. It contains several settings:

- Enable fabric snapshot creation:** A checkbox that is currently unchecked. Below it, a note states: 'Advanced server settings may be used by TAC for debugging purposes and not normally required for regular operation.'
- Display advanced settings and options for TAC support:** A checkbox that is currently unchecked. Below it, a note states: 'Splunk allows for a comprehensive view of application performance and infrastructure health.'
- Enable Splunk:** A checkbox that is checked. To its right is a help icon.
- Upload or paste Splunk license key:** A section with the instruction 'License files end with .license'. It contains an 'Upload' button with a file icon.
- Paste Splunk license key:** A text area containing the following XML license key:

```
</features>
<add_ons>
  <add_on name="itsi" type="app">
    <parameter key="size" value="10"/>
  </add_on>
</add_ons>
<sourcetypes/>
<guid>20521A48-9DF3-4198-927F-330B50574FC2</guid>
</payload>
</license>
```

At the bottom right of the page are 'Cancel' and 'Save' buttons.

Access: Splunk UI is available at <https://<mgmtIP>/splunk-ui> or through the Splunk card in the Analysis Hub.

User and role management:

- Local Nexus Dashboard users, passwords, and roles are synchronized with Splunk.
- Splunk **admin** is different from Nexus Dashboard **admin** user and access to Splunk **admin** user is not available. **nd_admin** is equivalent to Nexus Dashboard admin user.

Data collection features:

- After the embedded Splunk is enabled for Audit logs, Anomalies (System and Fabric), and Advisories (System and Fabrics) are automatically streamed to this Splunk instance. No other data is streamed to Splunk.
- The embedded Splunk configuration is different from the external Splunk configuration available on the **Admin > System Settings > Remote Streaming** page. You cannot customize the data streamed from Nexus Dashboard to the embedded Splunk instance.

Backup and restore:

- When the embedded Splunk configuration is enabled, the Nexus Dashboard backup includes both the configuration and any custom dashboards created or edited on the Splunk instance. Restoring an Nexus Dashboard backup will also restore the Splunk configurations and dashboards. Dashboards created and configuration changes by users are restored from the backup previously taken.
- Support for restoring the operational data collected within Splunk is not available.

Compliance rules

Compliance rules are created to generate anomalies where compliance can be violated or satisfied. Once you create compliance rules, you can generate the Compliance Report to check how much the fabrics and networks align to the rules.

Click **Manage > Anomaly and Compliance 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, Fabric, 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.

Interpretation of compliance rules

The following table lists some examples of compliance rules and what condition they create.

Compliance Rule	Condition Created
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 VRF instances where the tenant is "xyz" and the VRF instance name contains "c"	When an EPG under tenant "abc" that is in a VRF instance with DN uni/tn-xyz/ctx-abcde is selected, verify that both the tenant and the VRF instance criteria match. An EPG under tenant "abc" that is in a VRF instance with DN uni/tn-xyz1/ctx-abcde is not selected because the VRF instance 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 instance with DN uni/tn-rrr/ctx-sss	An EPG under tenant "abc" that is in a VRF instance with DN uni/tn-rrr/ctx-sss is selected because the VRF instance tenant matches.

Compliance analysis

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 **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 fabrics or individual view for a specific fabric. 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 attributes like 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, and 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 attributes like Name, Rule Type, Enforcement Status, and Verified. The **Create Compliance Rule** button takes you to the rule creation page.

Compliance anomalies

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

The number of anomalies raised is defined by the number of rules associated with a snapshot. For example, if an assurance group runs a compliance analysis on a snapshot every 15 minutes, and there are two rules associated with the snapshot, two anomalies will be raised.

Guidelines and limitations for compliance

Guidelines for compliance

- A single compliance rule can be associated with multiple fabrics.
- You can have a maximum of 30 active Communication Compliance rules and 600 active Configuration Compliance rules per fabric. 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 fabrics, it is recommended that you do not start a bug scan for those fabrics.
- Fabric list can be modified at any point in time.
- Name of the rule is unique across fabrics.
- 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 for creating a compliance rule

- When you create a compliance rule, you can add a custom description, which appears in the compliance violation anomaly.
- Compliance Rules are created at the fabric level.
- A compliance rule can either be offline or online.

Guidelines for import configuration compliance

- 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.

Guidelines for compliance rule naming

- Name should be a minimum of three characters
- Name should not include special characters
- Name should be unique.
- No two rules can have the same name.

Verified scalability guidelines 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 $5 \times 15,000$ 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 fabric.

Guidelines for template-based compliance

- 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 white spaces. Pretty JSON files will have white spaces to support indentation. To reduce the file size, you can remove white spaces and upload the file.
- 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.

- 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.

Limitations for compliance analysis

- No telemetry is available for offline analysis.
- The compliance report is generated once every two hours.

Limitations for Integrating Splunk for advanced monitoring and compliance

- Splunk is only supported on 1-node physical clusters with a ND-NODE-G5S (UCS-C225-M8) node type.
- Splunk is not supported as a cluster and is supported only as a standalone instance.
- Only locally-created users can log into Splunk.

Create compliance communication rule

1. Provide the name and description for your rule.
2. Select the fabrics you would like to apply the rule to.

You can pick one, or many, or all fabrics.

3. Choose to enable or disable the rule, then click **Next**.
4. In the **Compliance Rule Type** field, choose **Communication**.
5. 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**. The communication types are applied between two different object groups.
6. In the **Object Type** fields and the **Traffic Selector** area, choose the appropriate objects and traffic selector.
7. Select the appropriate criteria for both groups. Select any object type and the corresponding matching criteria object. See [Matching criteria](#) for the available object types and to understand how the various matching criteria objects can be defined.
8. 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. See [Communication compliance](#) for more information about the communication types and the traffic selectors.
9. After you complete defining the objects, criteria, traffic restrictions as appropriate for your fabric/s, you can view the entire overview of the rule create and click **Save Rule** to complete the configuration.
10. 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.



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

column.

Create compliance rule with snapshot selection

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**.



When using offline datasets to save compliance rules, ensure the offline fabrics are configured with Premier. Offline datasets configured with Essentials result in an incorrect UUID value in the payload, preventing the rules from being saved. This issue is observed in Nexus Dashboard 4.3.1 and NDI 6.7.0 releases.

4. In **New Rule**, click **Save**. Nexus Dashboard starts performing a check.
5. To download the snapshot, click the **Download** link from **Settings**.

Import configuration compliance

1. Under **Compliance Rule Type**, choose **Configuration**.
2. In the **Base Configuration Settings** field, choose **Import Configuration**. 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**.
3. Drag and drop your file into the provided field to upload. Click **Save**.

Create compliance rule with manual configuration

1. Provide the name and description for your rule. You can choose to Enable or Disable state.
2. Select the fabrics you would like to apply the rule to. You can pick one, or many, or all fabrics.
3. In the **Compliance Rule Type** field, choose **Configuration**.
4. In the **Base Configuration Settings** field, choose **Manual Configuration**.
5. 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. Select any object type and the corresponding matching criteria object. See [Matching criteria](#) for the available object types and to understand how the various matching criteria objects can be defined. See [Manual configuration compliance](#) for information about the attribute requirements.
6. 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.

7. You can view the entire overview of the rule you want to create and click **Save Rule**. Nexus Dashboard to start performing a check based on the Naming compliance requirements that you specified.

- 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.



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. See [Manual configuration compliance](#).

Create template-based compliance

- In the **Base Configuration Settings** field, choose **Template Based Compliance**.
- 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.
- Click **Save**.

For more information about template syntax, see [Templates to configure object selectors](#). For information on how to configure object selectors for the template, see [Templates to configure object selectors](#).

Trigger a compliance analysis

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

- Navigate to **Analyze > Analysis Hub > Compliance**.
- Select a fabric from the dropdown menu.
- Select the date for which you would like to see the report.

Templates to configure object selectors

When you create a configuration rule using manual configuration, only a few specific object selectors are supported (such as BD, EPG, VRF). By using a template, you can select any object and apply match criteria on its attributes.

An object can be any managed object from Cisco 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. You can configure object selectors for selection and match criteria and based on tags and annotations.

Selection, status, and match criteria

For naming compliance, the compliance rules are on the name and nameAlias fields that are indicated by **MATCH**.

- STATUS** defines the state of the specific object in the template, whether an object exists or does not exist. The **STATUS** criteria can be defined using one of the following keywords.

- MUST_EXIST
- MUST_NOT_EXIST

The following is a syntax example:

```
{
  "vzBrCP":
  {
    "attributes":
    {
      "STATUS": "<status selected>",
      "SELECT(name)": "<KEY_WORD>(<value>)",
      "MATCH(nameAlias/name)": "<KEY_WORD>(<value>)"
    }
  }
}
```

- The **SELECT** and **MATCH** criteria can be defined using one of the following keywords. The **MATCH** criteria is used to define the Compliance rule. **SELECT** allows to define a criteria to select group of objects and **MATCH** allows to define attributes and values that those selected objects must have. These compliance rules will be applied on objects that are selected using the **SELECT** criteria.

- STARTS_WITH
- ENDS_WITH
- EXACT
- OR
- REGEX

Syntax for SELECT:

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

Syntax for MATCH:

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



Attribute_name can be 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:

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

```

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

```

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

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**.

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)"
    }
  }
}

```

```
}
```

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".

▪ 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>)"

                }
            }
        }
    ]
}

```

- Template With Attribute Value **NULL** or **EMPTY**

The following are examples of templates where the attribute value is null or empty.

```

"REGEX(^.{0}$)"
"EXACT()"
"OR(test, )" ← use space

```

```

{
    "fvTenant":
    {
        "attributes":
        {
            "MATCH(annotation)": " OR(orchestrator:msc, )",
            "SELECT(name)": " REGEX(aepg_aepg_imd_tnt_pass_[0-9]+)",
        }
    }
}

```

```
}
```

For the procedure to configure Object Selectors for Naming Compliance using the above template, see [Create template-based compliance](#).

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.

*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>)"
          }
        }
      }
    ]
  }
}
```

*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.

*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 [Create template-based compliance](#).



When using the steps to [Create 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.

Communication compliance

Communication types

- **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 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.

Communication type and traffic selector rules selections with 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	• EPG • Tenant	• 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

Traffic selector rules available

Ether Type	Protocol Type
ARP	-
FCOE	-
IP	• All • EGP • EIGRP • ICMP • ICMPV6 • IGMP • IGP • L2TP • OJPFIGP • PIM • TCP • UDP
MAC_SECURITY	-
MPLS_UNICAST	-
TRILL	-

Manual configuration compliance

Attribute requirement that can be set according to the objects selected

Object	Associated Attributes
EPG	The associated attributes are: ▪ Preferred Group Member- The preferred group member can be equal to or not equal to either <i>Include</i> or <i>Exclude</i>. ▪ Infra EPG Isolation- The Infra EPG Isolation can be equal to or not equal to Unenforced/Enforced. ▪ QoS Class- The QoS Class can be equal to or not equal to Unspecified/Level 1/Level 2/Level 3.
VRF	The associated attributes are: ▪ Enforcement Preference- The enforcement preference can be set to equal to or not equal to Unenforced/Enforced. ▪ Enforcement Direction - The enforcement direction can be set to equal to or not equal to Ingress/Egress. ▪ Preferred Group - The preferred group can be set to equal to or not equal to Disabled/Enabled. ▪ BD Enforcement - The BD enforcement can be set to equal to or not equal to Yes/No.

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

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

Objects available as matching criteria for a selected object type

Object Type	Matching Criteria Object
EPG	▪ Tenant ▪ VRF ▪ BD ▪ EPG ▪ App profile ▪ L3 Out ▪ L3 InstP ▪ L2 Out ▪ L2 InstP
Tenant	▪ Tenant
BD	▪ Tenant ▪ VRF ▪ BD
VRF	▪ Tenant ▪ VRF
Contract	▪ Tenant ▪ Contract
Subject	▪ Tenant ▪ Subject
Filter	▪ Tenant ▪ Subject ▪ Filter

Define matching criteria objects

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>

Matching Criteria Object Type 1	How to define
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> flt - operator <i>value</i>



operator and *value* can be set to anything.

Operators for 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.
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.

Understanding the hardware and software lifecycles for your network using Conformance

Conformance reports enable 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 fabric for hardware and software conformance and weekly for each fabric 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 fabrics.

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 fabrics.

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 fabrics.

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

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 fabrics 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**.

Choose a fabric from the drop-down list.

OR

Navigate to **Manage > Fabrics**.

Choose a fabric.

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

Click **View Report**.

View conformance report



You can save conformance report as a PDF with the browser print option (Only supported on Chrome and Firefox).

1. Navigate to a Conformance Report. See [Access conformance report](#).
2. Choose a fabric or **All Fabrics** 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 fabrics 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, fabric 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, fabric 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, fabric, 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 fabrics summary, scale conformance, and scale metrics.

a. The All fabrics 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 fabrics and switches. The All Scale Metrics section displays if you choose **All Fabrics** from the drop-down menu.

i. The details for fabrics include fabric name, type, software version, controller metrics conformance, switch metrics conformance. Click the fabric name to view additional details.

ii. The details for switches include switch name, fabric 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 fabric, type, software version.

iv. Use the gear icon to customize the columns in the table.

e. The Fabric Level Scale Metrics and Switch Level Scale Metrics displays the scale metrics details for a fabric and switches associated with the fabric. These sections are displayed, if you choose one fabric from the drop-down menu.

i. The details for a fabric include metric, conformance status, and resource usage,

ii. The details for switches include switch name, fabric 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.

Analyzing endpoint flows using Connectivity Analysis

Connectivity Analysis allows you to analyze flows between two different endpoints, provide insight into how your endpoints are connected, and helps you spot where problems might be occurring. Starting with this release, connectivity analysis now includes multi-fabric capabilities.

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 in troubleshooting by narrowing down the root cause of the issue, including running forwarding path checks, software and hardware state programming inconsistencies through consistency-checks, and further details related to packets walkthrough.



Connectivity analysis supports microsegmentation endpoint groups (uSeg EPGs). You can trace forwarding paths and identify nodes causing drops for flows that include uSeg endpoint groups (EPGs). This feature validates policy-based redirect (PBR) traffic redirection.

Connectivity Analysis support for Layer 4 to Layer 7 services

Connectivity analysis (CA) supports enhanced Layer 4 to Layer 7 (L4-L7) services. This feature allows you to analyze and troubleshoot traffic flows that traverse L4-L7 devices in fully managed VXLAN EVPN fabrics. Connectivity analysis provides visibility into L4-L7 deployment scenarios to help you verify traffic paths through network services.

Connectivity analysis supports these L4-L7 use cases:

- Redirect to Service Chain - Traffic can be redirected through a sequence of Layer 4 to Layer 7 (L4-L7) service nodes (for example, firewalls, load balancers, or traffic optimizers) using enhanced Policy-Based Routing (ePBR). Connectivity analysis does not support service chaining. This redirection supports both intra-VRF (within the same VRF) and inter-VRF (between different VRFs) scenarios.
- Perimeter Service - Analyze traffic flows routed through security or network services positioned at the network perimeter. It facilitates traffic movement between different VRFs and Forwarding VRFs, ensuring segmentation and secure inter-VRF connectivity.
- Route Peering Service - This establishes dynamic routing adjacencies between L4-L7 service devices and network routers, allowing for flexible traffic steering and optimized service insertion within a single VRF or across multiple VRFs.
- Legacy Service Redirection - This utilizes traditional PBR to redirect traffic to legacy L4-L7 services, supporting both intra-VRF and inter-VRF deployments.

Explorer integration with connectivity analysis

The integration between explorer and connectivity analysis includes these features:

- Explorer to connectivity analysis integration: When you run an Explorer query between two endpoints, the system displays a connectivity analysis (CA) preview that shows the predicted topology.
- Full analysis: You can start a full connectivity analysis job from the Explorer results. The Previous Analyses panel displays links to the CA job records.
- Real-time job enrichment: If both the source and destination have associated endpoint security groups (ESGs) or group policy options (GPOs) during a CA job, the system performs an automatic Explorer check. The connectivity analysis report displays policy names and selectors on the starting switches. When the connectivity analysis job completes, the **Security** tab displays the results.

[nd ca explore integration] | [nd-ca-explore-integration.jpg](#)

Silicon One Packet Tracer integration

Connectivity Analysis now leverages the Silicon One Packet Tracer tool to capture and inspect live traffic for a given flow between two endpoints. This report's results are displayed in the same format as Embedded Logic Analyzer Module (ELAM) within the Connectivity Analysis, under **Analysis details > Data plane**. With the integration of L4-L7 services and Packet Tracer, the Connectivity Analysis topology now visually includes the service leaf and L4-L7 devices in the traffic path. For Cisco Silicon One switches, the packet report displays packet data for the receive (RxPP) and transmit (TxPP) post-processing pipelines.

Prerequisites for L4-L7 Connectivity Analysis

Before initiating any L4-L7 Connectivity Analysis flows, ensure the following conditions are met:

- Ensure the fabric has a Premier license and is in managed mode. Run the most recent fabric agents for connectivity analysis. Enable telemetry for the fabric to support end-to-end path visualization and data plane analysis. Connectivity analysis uses telemetry data to trace flows through Layer 4 to Layer 7 (L4-L7) service nodes.
- All L4-L7 services must be correctly configured at either the fabric or fabric group level (depending on your deployment use-case) within Nexus Dashboard. For more information see [Working with L4-L7 services](#)
- All participating fabrics must be fully managed VXLAN EVPN fabrics.
- Ensure that all switches run supported and officially released Cisco NX-OS versions.
- Update fabric agents to the most recent version before starting connectivity analysis (CA) jobs.
- Ensure all L4-L7 insertions and functions are correctly deployed. Perform recalc and deploy operations in Nexus Dashboard to ensure the configuration status is fully synchronized across the fabric.

Guidelines and limitations for Connectivity Analysis

- The NX-OS version requirements for Layer 4 to Layer 7 (L4-L7) support is NX-OS (9.3(7a) through 10.6(2) and later) and Silicon One Packet Tracer is supported only for Silicon One switches starting (10.5(2) and later).
- Nexus Dashboard 4.2.1 does not support these service features:
 - Service chaining, where flows traverse multiple firewalls in sequence.
 - Service devices that require network address translation (NAT), such as load balancers.
- Firewall is the only supported Layer 4 through Layer 7 (L4-L7) device type in Nexus Dashboard 4.2.1.
- Connectivity analysis for Layer 4 to Layer 7 (L4-L7) services supports fully managed Virtual Extensible LAN (VXLAN) EVPN fabrics only. Nexus Dashboard must manage the fabric and L4-L7 configurations to ensure path visibility. The system does not support L4-L7 configuration paths or reliable path stitching for unmanaged fabrics or fabrics in monitor mode.
- The Legacy Service Redirection use-case, which utilizes traditional Policy-Based Routing (PBR), is currently supported only on single-site deployments as per current L4-L7 limitations. For more information, see [Guidelines and limitations for Layer 4 to Layer 7 services](#), and [Layer 4 to Layer 7 Services Configuration](#).
- Service as Default Gateway - Connectivity analysis (CA) does not support the service as default gateway use case. In this configuration, Layer 4 to Layer 7 (L4-L7) service devices such as firewalls function as the default gateway for endpoints in the same VRF. Because the default gateway is on a firewall external to the fabric, Nexus Dashboard does not know the host and gateway endpoints. Security administrators manage the gateway MAC and IP addresses on the firewall outside of Nexus Dashboard configurations. Nexus Dashboard and the switches cannot derive these parameters to reconstruct the traffic path.



If a Layer 4 to Layer 7 (L4-L7) service device, such as an external firewall, functions as the default gateway, Nexus Dashboard might not know the host and gateway endpoints. If you troubleshoot a flow between endpoints in different subnets using MAC addresses, the connectivity analysis job fails. In this scenario, the system displays a message that instructs you to use IP addresses instead of MAC addresses. You cannot start an IP-based troubleshooting job if Nexus Dashboard does not know or discover the specific IP addresses.

- Packet tracer does not provide a specific drop reason code or description, such as an MTU mismatch or an access control list (ACL) drop. However, the system identifies the point of drop by analyzing these two capture points on Silicon One switches:
 - from the external network to the network processing unit (NPU), and
 - from the NPU to the external network.

Connectivity Analysis options

- Embedded Logic Analyzer Module (ELAM)--ELAM is a diagnostic tool that helps troubleshoot ethernet traffic flows. It captures the packet from an active flow and analyzes the ethernet frames for packet drops. ELAM requires an active flow between the source and destination hosts. You can enable this option to analyze an available active flow.

- Operations, Administration, and Maintenance (OAM)--For standalone NX-OS fabrics, OAM is a protocol for monitoring and troubleshooting ethernet networks. You can enable this option to locate potential drops for active hosts or to track details such as reachability and probable route of the flow in a VXLAN EVPN-based fabric topology, without the need of active traffic between the hosts. OAM is supported only on VXLAN fabrics.



The OAM option is not available for NX-OS multi-fabric environments.

- Consistency Checker--For standalone NX-OS fabrics, Consistency Checker ensures system consistency and aids in root cause analysis and fault isolation by verifying the alignment between software and hardware tables. These checks are performed within each switch between data plane and control plane for all the networking entities related to the chosen endpoint conversation. You can enable this option to detect control plane and data plane configuration and operational inconsistencies along the flow between specified endpoints or routes.

The checks performed by Connectivity Analysis include:

- Topology checks such as overall health, connectivity of leaf switches, spine switches, or remote leaf switches
- VRF instance and bridge domain mappings for endpoints
- Interfaces connectivity such as PC, vPC, SVI, Breakouts, and Sublfs
- Routing tables, EPM, and EPMC tables
- L3Out information and mapping
- Adjacency (ARP) tables
- Tunnel information
- Displays Council of Oracle Protocol (COOP) synthetic route tables on spine switches for ACI fabrics only

Layer 2 ToR support

For standalone NX-OS fabrics, with Layer 2 ToR support, Connectivity Analysis offers the following capabilities:

- Incorporate the device into the topology of a connectivity analysis job.
- Provide detailed node-level flow information, including ingress interfaces and egress paths.
- Initiate ELAM (Embedded Logic Analyzer Module) and capture packet details on ToR switches.
- Conduct consistency-check validations on ToR switches.

Supported topologies

- ToR switch with port channel directly connected to a leaf switch.
- ToR switch connected to leaf switches in a vPC pair.
- ToR switches with port channels connected to leaf switches individually. The leaf switches are in a vPC pair.
- ToR switches in vPC with leaf switches along with ToR switches connected to hosts in a vPC pair.

Guidelines and limitations (connectivity analysis): For all fabrics

- You can submit up to 10 jobs.
- At any point of time, you can run only 1 connectivity analysis job per fabric. You can stop a job in the queue and run another job.
- You can run Connectivity analysis only on online fabrics.
- ELAM captures the VXLAN outer packet in multi-site fabrics for both ACI and NX-OS.
- ELAM capture is not a part of the connectivity analysis process, hence it does not provide packet captures during connectivity analysis.
- Nexus Dashboard supports only straight through FEXs. This support applies to connectivity analysis, flow telemetry, traffic analytics, and endpoint analytics.
- Egress paths from transit fabrics to border gateway groups may not all appear green. This happens because ELAM can only capture a single packet at a time, showing just one active path. However, if ELAM is triggered at the border, an active path exists even if some paths from transit fabric to border switches appear grey.
- Troubleshooting multi-fabric flows between ACI and standalone NX-OS fabrics is not supported.
- Cisco Silicon One Packet Tracer supports only Cisco N9336C-SE1 and later switches.

Guidelines and limitations (connectivity analysis): For fabric groups

Only full-mesh topologies are supported.

Guidelines and limitations (connectivity analysis): For Cisco ACI fabrics

- Connectivity analysis feature is supported on Cisco APIC release 6.0(2h) and Cisco ACI-Mode Switch release 16.0(2h) and later.



If your ACI fabrics are running version of 6.1.2 or later and your Nexus Dashboard version is 3.2.1 or earlier, you will not be able to perform connectivity analysis.

- Nexus Dashboard Connectivity Analysis provides limited support for Layer 4 to Layer 7 (L4-L7) devices in Cisco ACI fabrics. The following conditions and limitations apply:
 - Only a single firewall device is supported, Service chaining is not supported.
 - The firewall device must be attached in 2-arm mode to a single leaf.
 - Firewall arms must be attached via Bridge Domains (BD) using Policy-Based Redirect (PBR).
 - Connectivity Analysis is performed between endpoints in the fabric.
 - L4-L7 Service Graph is configured with Route Redirect.
 - Endpoints that are classified within an L3Out EPG are not supported.

- **Other** type of L4-L7 devices are not supported.

Guidelines and limitations (connectivity analysis): For standalone NX-OS fabrics

- Connectivity analysis is supported on Cisco NX-OS release 9.3(7a) and later.
- Connectivity analysis job is not triggered if all onboarded devices are shown as incompatible.
- Update the fabric agent on the switches to the latest version. Updating the fabric agent does not affect the traffic, so you do not need to reload the switch. The **View Nodes** banner displays the number of compatible and up-to-date nodes. Click **Update** to initiate the update process. After the update is complete, on the **Node** page, click **Refresh** to view the status.

- You can update fabric agent packages only when all the following requirements are met:
 - Telemetry is enabled for the fabric.
 - The fabric uses the Premier license tier.
 - Metadata is available through Cisco Intersight Cloud or an uploaded metadata bundle.
 - The fabric is in managed mode, or the fabric is in monitor mode with **Image management in monitor mode** option enabled.
- OAM supports single-site VXLAN fabrics only.
- OAM is only supported between VTEPS and as a result the OAM path will be displayed between Layer 3 networks.

Supported topologies

- Endpoint combinations:
 - EP-EP
 - EP-L3Out
 - L3Out-EP
 - L3Out-L3Out
- Conversation types:
 - L2, L3, L4 (ICMP/TCP/UDP)
 - ACI and NX-OS fabrics do not support ICMPv6
 - V4 and V6 support
 - Transit and Proxy flows
 - Shared Service
- Topologies for Cisco ACI fabrics:

- Single-Pod and Multi-Pod
- ACI Multi-Site
- Remote Leaf-Direct
- M-Topology (stretched fabric design)
- vPC
- 3-tier architectures
- Topologies for standalone NX-OS fabrics:
 - VXLAN
 - VXLAN Multi-Site
 - vPC
 - Classic LAN

Create connectivity analysis

Follow these steps to create a connectivity analysis.

1. Navigate to **Analyze > Analysis Hub > Connectivity**.
2. Click **New Analysis**.

← Connectivity Analysis

Create Connectivity Analysis

Fabric Agent Status ⓘ [View all](#)

3 Compatible ✔

Analysis details [Run analysis](#)

Fill out a few details and select which additional analysis you'd like to perform

Source

IP address *

Protocol

Port

☒ MAC

Destination

IP address *

Port

Options

Certain additional analyses may take extra time to complete. [Learn more](#)

☒ **Active path**
Adds ~30s

In the **New Analysis** page, you can view the **Fabric agent status** card which shows the number of onboarded fabrics and their compatibility status.

3. Click **View all**.


The **Fabric Agent Status** dialog box opens. You can view all the onboarded fabrics and compatibility status of the fabric with respect to a specific fabric agent that runs across all the NX-OS switches.



If all the fabrics are compatible, the **Status** card in the **Create Connectivity Analysis** page shows status as **Compatible**.

4. Click **Update Agent** to update the fabric agent for each fabric.

You can update fabric agents for each fabric or for all the fabrics.



You can click a fabric to view the switches that are part of the fabric and their agent versions in the **Switches for *fabric_name*** dialog box.

For standalone NX-OS fabrics, the fabric agent is installed on each switch as an RPM package. The **Current CA Version** column shows the fabric agent package version installed on the switch. The **Latest CA Version** column shows the package version available in the metadata. You can retrieve metadata from Cisco Intersight Cloud or upload it manually.

You can update fabric agents on a fabric in monitor-mode only when all the following requirements are met:

- o Telemetry is enabled for the fabric.
- o The fabric uses the Premier license tier.
- o Metadata is available from Cisco Intersight Cloud or an uploaded bundle.
- o Image management in monitor mode is enabled for the fabric.

If a fabric agent package update fails, review the update details for the switch. An update can fail if the package is unavailable, the package cannot be copied or installed on the switch, the switch does not have sufficient bootflash space, or in-band reachability to the switch is unavailable.

5. Complete the following for Layer2 and Layer 3 parameters.

- a. From the **IP address** drop-down list choose **IP address** or **MAC address** for **Source** and **Destination** to analyze the flow between two endpoints.

In previous releases, you could only run connectivity analysis jobs for NX-OS fabrics in a fabric group. Beginning with Nexus Dashboard release 4.3.1, you can now run connectivity analysis jobs to analyze the flow between two endpoints where the endpoints can be in ACI or NX-OS fabrics in a fabric group.

In addition, you can also run connectivity analysis at the fabric group level by navigating to that fabric group's **Overview** page and clicking **Actions > Start connectivity analysis job**. When you navigate to the **Connectivity analysis** page this way, the endpoints available for the connectivity analysis are filtered based on this fabric group.

- b. Choose the source endpoint from the drop-down list or enter the endpoint. A maximum of 20 IP or MAC addresses are displayed at a given time.

- c. You can also manually populate the Layer2 and Layer 3 parameters. Click **Enter Details Manually** to enter the source IP or MAC address, destination IP or MAC address, and fabric type. For ACI fabrics, also enter the source tenant, source VRF instance, destination tenant, and destination VRF instance. For standalone NX-OS fabrics, also enter the VRF instance and source VLAN.
- d. Choose the destination endpoint from the drop-down list or enter the destination endpoint.
6. For standalone NX-OS fabrics, choose **VXLAN** or **Classic** as the fabric type.
7. Complete the following for Layer 4 parameters.
 - a. From the Protocol drop-down menu, choose **ICMP**, **TCP**, or **UDP** protocol.
 - b. Enter the source and destination port number.
8. Choose the **Analysis Options**.
 - a. Check **ELAM** option to analyze an available active flow.
 - b. For standalone NX-OS fabrics, check **OAM** option to locate potential drops for active hosts or to track details such as reachability and actual route of the flow in a VXLAN EVPN based fabric topology, without the need of active traffic between the hosts. OAM is supported only on VXLAN fabrics.



The OAM option is not available for NX-OS multi-fabric environments.

- c. Check **Consistency Checker** option to detect control plane and data plane configuration and operational inconsistencies along the flow between specified endpoints or routes.



You cannot select both ELAM and OAM options for Connectivity Analysis.

9. Click **Run Analysis**.
10. After the Connectivity Analysis is completed, the analysis is displayed in the **Connectivity Analysis Jobs** table. Navigate to **Analyze > Analysis Hub > Connectivity** to view the Connectivity Analysis Jobs. The Analysis is assigned a default name and you can rename the analysis.
 - a. Select the analysis and then from the Actions drop-down menu click **Rename Analysis** to rename.

OR

 - b. Click on analysis name. In the **View Connectivity Analysis** page, from the Actions drop-down menu click **Rename Analysis** to rename.

View connectivity analysis

Follow these steps to view connectivity analysis.

1. Navigate to **Analyze > Analysis Hub > Connectivity**. The Connectivity Analysis jobs are displayed.

← Analysis hub

Connectivity Analysis

Analyze flows between two different endpoints, get insight into how your endpoints are connected, and spot where problems might be occurring. [Learn more](#)

Refresh New analysis

0 Running 0 Queued 1 Complete 0 Failed

Saved analyses Current

Filter by attributes Actions

Name	Job status	Flow status	Job ID	Source fabric	Destination fabric	Source IP	Destination IP	Analysis options	Start time	End time
NX-SITE1 NX-SITE2 02-02-2026 18:37:34	Completed	Success	FSV-23b4d9c8-0038-11f1-b73e-624e1d95ad92 NX-SITE1	NX-SITE1	NX-SITE2			Active path	Feb 02 2026 06:37:34 PM	Feb 02 2026 06:39:41 PM

- Choose a time range from the drop-down menu.
- The Summary area displays the overall status of the Connectivity Analysis jobs and the flow status.
- Use the filter bar to filter the list of jobs. The Connectivity Analysis table displays filtered jobs.
 - Click the column heading to sort the jobs in the table.
 - Click the gear icon to configure the columns in the table.
 - Hover around a failed Flow Status to learn more.
- Click **Name** to view the Connectivity Analysis details. The View Connectivity Analysis page displays the input parameters you had entered for the job, the job details, and topology.

Analysis details

Analysis completed successfully on Jan 30 2026, 08:11:52 AM

Re-run analysis

Fabric NX-SITE1	Source IP address	Source VRF vrf_50001	Source Tenant default	Source port number N/A	Start time Jan 30 2026, 08:09:33 AM
Fabric type NXOS	Destination IP address	Destination VRF vrf_50002	Destination Tenant default	Destination port number N/A	End time Jan 30 2026, 08:11:52 AM
				Protocol N/A	Run time 2 Minutes 10 Seconds
				Selected analysis options Active path	

Overview **Completed** Security Unavailable

Path visualization

The diagram illustrates the network topology for the connectivity analysis. It shows a Source node connected to a Destination node via a path that includes mns1-leaf3, mns1-leaf1, mns1-leaf2, and a Firewall. The path also connects to the ISN-EXT External Fabric and the NX-SITE2 NDFC Fabric. The diagram is titled 'Path visualization' and includes a 'Routing' sidebar on the right.

- In the **Source** field, enter the source site name/VRF name/source IP.
- In the **Destination** field, enter the destination site name/VRF name/destination IP.
- Click **Show Job Details** to view the job details.

- d. Click **Re-run Analysis** to run the analysis again.
- e. Below the **Show Job Details**, you can see a banner that displays the status of the job. A green banner represents a successful analysis and a red banner represents a failed analysis.
- f. You can view the routing status between the source end point and destination end point in the **Routing** card. Routing card states **OK** or **Not OK**.




You will see a green check mark, if the routing status is okay and a red cross (x) mark if there are any routing errors. You can click the red cross (x) mark to open the **Event Log** page. The routing errors will be highlighted in the **Event Log** page.

- g. For standalone NX-OS fabrics, use toggle enable or disable **Highlight Active Path**. When you enable **Highlight Active Path**, all the OAM paths in the topology are highlighted.
- h. Click **Re-run Analysis** to run the analysis again.
- i. In the topology area, you can visualize hierarchical view of the fabric. You can double-click on the node to view interconnections of the nodes in the fabric. The active path between nodes is highlighted in green color. See [Topology](#).



Starting from Nexus Dashboard release 4.0.1, connectivity analysis includes multi-fabric capabilities.

Path visualization



- j. Click a node to view the tooltip. The tooltip displays the node name, node type, and the ingress and egress connections for that node, as well as OAM information if applicable. In the ingress and egress connections, only physical interfaces are displayed. For standalone NX-OS fabrics, egress information is not displayed on the first VTEP node.
- k. Click **Analysis Details** to view the path and data plane information.
 - i. Click **Paths** to view path details such as ingress and egress information, and OAM information if applicable. In the ingress and egress connections area logical interfaces are displayed.

Analysis Results for scaleleaf-204



Paths Data Plane

General

Source IP

Destination IP

Ingress and Egress

Ingress Information

Local Ingress Logical Interface

Local Ingress Physical Interface

-

eth1/51

✓ Active Path

Egress Information

Local Egress Logical Interface

Local Egress Physical Interface

Peer Device

Peer Ingress Physical Interface

-

eth1/1

Destination

Destination

✓ Active Path

- ii. Click **Data Plane** to view the analysis options results of **External > NPU** and **NPU > External** page.

- **External > NPU**: Displays the packet as it enters the network processing unit (NPU) from the external network.

← Back

Analysis results for trex-nx-dansberg2

Paths Data Plane

External → NPU NPU → External

Inner L2 Header

Destination MAC	04:e3:
Source MAC	80:27:

Inner L3 Header

Destination IP	10.
Don't Fragment Bit	set
IP CheckSum	0x4cae
IP Packet Length	156
IP Version	4
Source IP	10.

Inner L4 Header

Checksum	0x0000
Destination Port	4789
Packet Length	136
Source Port	46605

Outer L2 Header

Destination MAC	04:e3:
Source MAC	80:27:

Outer L3 Header

- **NPU > External:** Displays the packet as it exits the NPU toward the external network.

← Back

Analysis results for trex-nx-dansberg2

Paths

Data Plane

External → NPU

NPU → External

Inner L2 Header	
Destination MAC	04:e3:
Source MAC	80:27:
Inner L3 Header	
Destination IP	10.:
Don't Fragment Bit	set
IP CheckSum	0×4cae
IP Packet Length	156
IP Version	4
Source IP	10.:
Inner L4 Header	
Checksum	0×0000
Destination Port	4789
Packet Length	136
Source Port	46605
Outer L2 Header	
Destination MAC	04:e3:
Source MAC	80:27:
Outer L3 Header	

iii. Click **ELAM** to view the ELAM report. Click **View Full Report** to download the report.

Analysis Results for scaleleaf-204



Paths **Data Plane**

Data Plane Details

ELAM

Basic Information

Device Type	LEAF
Packet Direction	egress
Incoming Interface	Eth1/51

Inner L2 Header

Inner Destination MAC	00:00:00:00:00:00
Source MAC	00:00:00:00:00:00
802.1Q tag is valid	no
CoS	0
Access Encap VLAN	0

Outer L2 Header

Destination MAC	00:00:00:00:00:00
Source MAC	00:00:00:00:00:00
802.1Q tag is valid	yes
CoS	0
Access Encap VLAN	2
VN-Tag is valid	no
Src VIF(in from leaf/IPN)	0
Dst VIF(out to leaf/IPN)	0

Inner L3 Header

L3 Type	IPv4
DSCP	0
Don't Fragment Bit	0x0
TTL	60
IP Protocol Number	ICMP
Destination IP	10.0.0.1
Source IP	10.0.0.2

Outer L3 Header

L3 Type	IPv4
DSCP	0
Don't Fragment Bit	0x0
TTL	31
IP Protocol Number	UDP
Destination IP	10.0.0.1
Source IP	10.0.0.2

Outer L4 Header

L4 Type	iVxLAN
Don't Learn Bit	1
Src Policy Applied Bit	0
Dst Policy Applied Bit	0
sclass(src pcTag)	0x4002
VRF or BD VNID	2818048(0x2B0000)

- iv. For standalone NX-OS fabrics, click **Consistency Checks** to view the consistency check results.

Analysis Results for topo-3-fx-leaf-1

Paths Data Plane

Data Plane Details

ELAM Consistency Checks

Ethernet1/43

✔ Spanning Tree Protocol state validator

show consistency-checker stp-state vlan 2401 brief

No Issues Found

✔ L2 Switchport state validator

show consistency-checker l2-state interface Ethernet1/43 brief

No Issues Found

Analysis Results for topo-3-fx-leaf-1

✔ Gateway mac state validator

show consistency-checker gwmacdb interface vlan 2401 brief

No Issues Found

✔ L3 physical routed port state validator

show consistency-checker l3-interface interface vlan 2401 brief

No Issues Found

✔ L2 MAC state validator

show consistency-checker vxlan l2 mac-address 0050.0100.0001 module 1 brief

No Issues Found

Analysis Results for topo-3-fx-leaf-1

✔ VxLAN VLAN state validator

show consistency-checker vxlan vlan 2401 brief

No Issues Found

✔ Physical Front Panel Port Link state validator

show consistency-checker link-state interface Ethernet1/43 brief

No Issues Found

loopback1

✔ Physical Front Panel Port Link state validator

show consistency-checker link-state interface Ethernet1/50 brief

No Issues Found

✔ Physical Front Panel Port Link state validator

show consistency-checker link-state interface Ethernet1/49 brief

No Issues Found

✔ L3 physical routed port state validator

show consistency-checker l3-interface interface Ethernet1/49 brief

No Issues Found

✔ L3 physical routed port state validator

show consistency-checker l3-interface interface Ethernet1/50 brief

No Issues Found

✔ L3 route state validator

show consistency-checker forwarding single-route ipv4 10.3.0.3/32 vrf default brief

No Issues Found

- v. For standalone NX-OS fabrics, click **OAM** to view the OAM report.
- l. Click **Node Details** to view the node details in inventory. See [Inventory](#).

Manage connectivity analysis

Follow these steps to manage connectivity analysis.

- 1. Navigate to **Analyze > Analysis Hub > Connectivity**.
- 2. Click **Name** to view Connectivity Analysis details.

Analysis > Analysis Hub > Connectivity Analysis > Topo3-Site1 Topo3-Site2 14-01-2025 15:15:04

View Connectivity Analysis

Source

Topo3-Site1 / vrf_50001 / 182.31.15

Destination

Topo3-Site2 / vrf_50001 / 182.31.15

Layer 4 Parameters

Protocol

Source Port Number

Destination Port Number

Analysis Options

☒ ELAM

☐ OAM

☐ Consistency Checker

Stop Analysis

Re-Run Analysis

Run Reverse Analysis

Rename Analysis

Show event log

Fabric Type

Classic

Show Job Details

Re-Run Analysis

View Connectivity Analysis

Source

Topo3-Site1 / vrf_50001 / 182.31.1.5

Destination Topo3-Site2 / vrf_50001 / 182.31

Protocol

Destination Port Number

Analysis Options  ☒ ELAM ☐ OAM ☐ Consistency Checker

Show Job Details Re-Run Analysis

3. From the Actions drop-down menu choose **Stop Analysis** to stop the analysis.
4. From the Actions drop-down menu choose **Re-Run Analysis** to run the analysis again.
5. From the Actions drop-down menu choose **Run Reverse Analysis** to run the analysis in the reverse direction.
6. From the Actions drop-down menu choose **Rename Analysis** to rename the analysis.
7. From the Actions drop-down menu choose **Show Event Log** to view the logs for the analysis. In the event log, you can view the error message for a failed analysis.

Operator	Description
<	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.

Monitoring the latency, congestion, and drops for your network, using traffic analytics

Traffic analytics enables you to monitor your network's latency, congestion, and drops. In Nexus Dashboard, you can now configure the traffic analytics mode best suited for your fabric type and requirements. You must enable traffic analytics at the cluster level before configuring it for your fabrics.

You can choose to configure one of these Traffic analytics modes.

- Traffic analytics full
- Traffic analytics compatibility
 - Sub-modes
 - (NetFlow)
 - (sFlow)
- Disabled



Nexus Dashboard determines the appropriate **Traffic analytics compatibility > Sub-mode (NetFlow or sFlow)** and automatically chooses the right **Sub-mode** for your fabrics.

Traffic analytics full

The **Traffic analytics full** mode matches well-known TCP and UDP (optional) Layer 4 ports to their corresponding service endpoint categories to automatically discover services running in your network. Nexus Dashboard then assesses service performance based on thresholds for these metrics.

- Latency—Measures the overall time in microseconds it takes a packet to go between the ingress and egress leaf switches for specific traffic flow. Latency is tracked for both ingress and egress traffic between a service endpoint and its clients.
- Congestion—Measures network bandwidth utilization and quality of service (QoS) activation mechanisms to determine if a service is experiencing network congestion. For ACI fabrics, this also measures priority flow control (PFC) and explicit congestion notification (ECN) counters.
- Drops—Measures the score or number of dropped packets versus transmitted packets considering factors such as CRC errors, faulty cables, and other devices.

An anomaly is raised if there is any deviation in the performance metrics such as latency, congestion, and drops. The performance score is calculated for each conversation and aggregated to the service endpoint or endpoint level to raise anomalies.

The performance score is calculated based on these factors.

- Congestion—Consistent congestion avoidance active between endpoints is calculated.
- Latency—Deviation from the average latency of the previous conversations is calculated.
- Drops—Directly correspond to an issue with the conversation or service.

Using traffic analytics, you can

- monitor traffic pervasively
- report performance issues using anomalies raised for performance metrics
- sort top talking services and clients and determine the top talkers in the system
- determine the SYN or RST counts per service, and
- troubleshoot conversations or flows on-demand.

You can configure the **Traffic analytics full** mode from the Nexus Dashboard **Fabric > Edit Fabric Settings** page. For more information, see [Configure traffic analytics flow collection](#).



Traffic analytics is supported on ACI version 6.1.1 or later and NX-OS version 10.4(2) and later.

Traffic analytics compatibility

The **Traffic analytics compatibility** modes allow you to configure traffic analytics for fabrics that do not support the full mode capability, such as fabrics running ACI or NX-OS software versions not compatible with the full mode feature. Traffic analytics compatibility mode is also intended to support other fabrics that are not full mode capable. Note that traffic analytics compatibility mode has limited capabilities compared to full mode, review the limitations noted here for more information.



- Traffic analytics compatibility mode supports Netflow and sFlow modes in Nexus Dashboard 4.3.1.
- You must manually configure NetFlow for both ACI and NX-OS fabrics. Nexus Dashboard will not configure the fabric to send NetFlow information. Similarly, you must manually configure sFlow for supported NX-OS fabrics.
- Traffic analytics compatibility mode currently supports only congestion anomalies. Latency and drop anomalies are not supported.
- View the **Traffic Analytics Capability** column on the **System Status** page to see which switches support **Traffic analytics full** mode and which support the **Traffic analytics compatibility** mode.

Enabling the traffic analytics compatibility mode helps with service identification using well known ports. You can configure the compatibility mode from the Nexus Dashboard **Fabric > Edit Fabric Settings** page. For more information, see [Configure traffic analytics flow collection](#).

Beginning with 4.3.1, the **Traffic Analytics compatibility** feature is enhanced to support NetFlow and sFlow based switches in your fabric. For more information, see [Traffic analytics support for Silicon One \(S1\) platforms](#).

Disabled

You can choose to configure the **Disabled** mode to turn off the traffic analytics feature for your fabric. When **Disabled** mode is configured, you cannot generate or view flow collection data for your fabrics. For more information on configuring the **Disabled** mode, see [Configure traffic analytics flow collection](#).

Traffic analytics conversations

A TCP conversation is a 4-tuple including source IP address, destination IP address, destination port, and protocol. A non-TCP conversation is a 3-tuple including source IP address, destination IP address, and protocol. In case a single client establishes multiple communication flows initiated by multiple source ports toward a service endpoint, all related statistics would be aggregated as a single entry in the Traffic analytics table. A service endpoint is defined by an IP address, a port, and a protocol.

For Cisco ACI fabrics, traffic analytics includes microsegmentation endpoint group (uSeg EPG) information. The EPG field in the Traffic Analytics table and conversation views displays the standard EPG name or the uSeg EPG name.

An anomaly is raised after the conversation rate limit is exceeded. Navigate to **Admin > System Settings > Flow Collection**. In the Traffic analytics status for the last hour area, you can view if the conversation rate approaches or exceeds the limits. You can also view if there are any Traffic analytics record drops.

Traffic analytics transit conversations

Beginning with 4.3.1, Nexus Dashboard provides more visibility into network conversations, particularly those involving endpoints external to the fabric. Visibility is now extended to conversations where both the source and destination endpoints are external to the fabric, such as traffic routed through an L3Out. In addition, the settings to enable or disable external conversation visibility is moved from **System Settings > Flow collections** to **Fabric Settings > Telemetry**.

Under telemetry fabric settings, a new option is added to Traffic analytics scope choices, **Inter-fabric + external** scope is added to the Traffic analytics configuration at the fabric level enabling more granular control over data collection. The **Inter-fabric + external** scope allows comprehensive traffic capture regardless of whether the client or service endpoint has been learned within the local fabric. This is also beneficial for Layer 2 (L2) traffic scenarios where traditional endpoint learning may not occur.

You can navigate to **Manage > Fabric > Edit fabric settings** and choose **Inter-fabric + external** from the **Traffic analytics scope** field to enable this scope for your fabric. When the **Inter-fabric + external** scope is enabled, external services are directly displayed on the **Traffic analytics services** page, eliminating the need to navigate to separate connectivity endpoint views to identify external service interactions. In addition to the **Inter-fabric + external**, the Traffic analytics scope field includes the **Inter-fabric** and **Intra-fabric** options. When you choose the Inter-fabric option, Traffic analytics monitors all traffic as long as the client or service endpoints are learned in the local fabric. Choosing the **Intra-fabric** option allows Traffic analytics to monitor traffic where both client and service endpoints are learned in the local fabric.

Upgrade behavior

The **Inter-fabric** scope also replaces the cluster-level multi-fabric enabled/disabled property for Traffic analytics. If you are upgrading previous versions, such as Nexus Dashboard 3.2 or 4.1, the transition is handled automatically during the upgrade process.

When you upgrade Nexus Dashboard, the multi-fabric behavior depends on the cluster-level multi-fabric setting:

Upgrading from Nexus Dashboard 3.2 to 4.1:

- The cluster-level multi-fabric property remains enabled, and all fabrics are set to inter-fabric scope.

Upgrading from Nexus Dashboard 4.1 to 4.2:

- If the cluster-level multi-fabric property is disabled, each fabric is set to inter-fabric scope.
- If the cluster-level multi-fabric property is enabled, each fabric is set to inter-fabric scope.



For greenfield deployments, **Inter-fabric + external** scope is not enabled by default after the upgrade. Each fabric is set to either intra-fabric or inter-fabric scope based on the existing configuration. If a fabric requires **Inter-fabric + external scope**, you must manually change the scope for that specific fabric.


Conversation tagging and statistics

Conversations involving remote (non-local) endpoints are now tagged for easier identification of Transit and L2/Transit traffic. The **L2/Transit** tag is applied to conversations between two remote (non-local) endpoints, or in L2-only Virtual Network Identifier (VNI) scenarios where endpoints are not learned. Similarly, **L2/Transit** tag is applied for L2-only VNI scenarios, providing explicit context. This classification helps in identifying traffic traversing the fabric without local endpoint learning.

- Flow enrichment: Flow records carry source and destination class (sclass and dclass) information. The system enriches these records to identify the endpoint security group (ESG) or microsegmentation endpoint group (uSeg EPG).
- API support: The EPG field in the Flow and Traffic Analytics API returns the EPG name or the uSeg EPG name based on the endpoint classification.

The following new columns are available in the **Analyze > Analyze Hub > Traffic analytics > Services** tab.

- **Location** – Displays the location of the Service Endpoint (SEP).
 - **Local** – Indicates learned locally in the selected fabric.
 - **External** – Indicates the endpoint is not local, fabric is displayed in the Fabric if endpoint is learned on a remote fabric managed by Nexus Dashboard, otherwise the Fabric is empty.
- **Conversation Attributes** – Aggregation of all conversation attributes seen from the clients.
 - **L4L7** – Indicates conversation is going through L4L7 device.
 - **L2/Transit** – Indicates both Client and Service endpoints are not-learned in Nexus Dashboard.
 - **External** – Indicates that some clients are non-local.
 - **Local** – Indicates that some clients are local.
- **Fabric** – Displays the name of the fabric where this service endpoint was originally learned. If no endpoint is learned, this column will be empty.



Enhancements to traffic analytics endpoints view

The **Endpoints** tab is now available on the Traffic analytics page to facilitate quick and easy navigation. The **Endpoints** tab aggregates all endpoints seen generating traffic in Traffic analytics. Both service and client endpoints are listed in this view.

In earlier releases, viewing client endpoint information required validating the endpoint at the fabric level. The newly introduced Endpoints view eliminates this requirement and provides a significant operational enhancement.

You can click **Analyze > Analyze Hub > Traffic analytics** to view the **Endpoints** tab along with the **Overview** and **Services** tabs. The Endpoints view is enhanced to help you identify the traffic details. The Endpoints table includes Conversation attributes, Location, Fabric, and Traffic analytics score columns. You can click on an Endpoint (IP address) to view the Traffic analytics details, such as services on the IP and other connections from the IP.

Traffic analytics for multi-fabric and external conversations

By default, traffic analytics monitors traffic that is external to the fabric in a multi-fabric configuration. Remote fabric endpoints are identified with the **External** tag and have a link that takes you to the respective endpoint page to assist with visibility.

To provide traffic analytics visibility for encapsulated VXLAN conversations flowing through a transit NX-OS fabric in a multi-fabric topology, configure the transit fabric as follows:

- Enable traffic analytics for the transit NX-OS fabric.

- Set the fabric **Traffic analytics scope** to **Inter-fabric + external**.
- Assign switch roles that support traffic analytics to all switches in the fabric. For example, use the **Aggregation** role for Classic LAN fabrics.
- Enable VXLAN EVPN overlay support on the switches by configuring the NX-OS CLI commands **feature nv overlay** and **nv overlay evpn**.

Traffic Analytics Full Inter-fabric + external I-ipn Current Refresh

Data is shown based on telemetry-monitored hardware. You can [Learn more about our methodology here](#).

Overview **Services** Endpoints

View **Service categories** by **Traffic analytics score**

Location == External x Fabric contains fab x Edit Clear all View endpoint details

Endpoint	Service port	Conversation attributes	Location	VRF	Switches	Interface	Fabric	Conversation Count (Total)	Traffic analytics score
20.16.144.1	4791	L2/Transit	External	myvrf_50004	n9k-leaf-3	eth1/1	-fab3	10	Major
20.1.11.1	443	External, L2/Transit	External	myvrf_50001	n9k-leaf-11	eth1/1	-fab1	10	Healthy
20.1.11.1	443	External, L2/Transit	External	myvrf_50001	n9k-leaf-11	eth1/1	-fab1	10	Healthy
20.1.13.1	443	External, L2/Transit	External	myvrf_50001	n9k-leaf-1	eth1/49	-fab1	10	Healthy
20.1.13.1	443	External, L2/Transit	External	myvrf_50001	n9k-leaf-1	eth1/49	-fab1	10	Healthy
20.11.133.2	53	L2/Transit	External	myvrf_50003	n9k-leaf-2	eth1/1	-fab2	5	Healthy
20.11.144.1	3389	L2/Transit	External	myvrf_50003	n9k-leaf-3	eth1/1	-fab3	10	Healthy
20.11.144.1	5678	L2/Transit	External	myvrf_50003	n9k-leaf-3	eth1/1	-fab3	5	Healthy

Traffic analytics support for traffic redirection to L4-L7 services

Nexus Dashboard traffic analytics capability now supports monitoring of traffic redirected to Layer 4 to Layer 7 (L4-L7) services. A new tag (L4-L7) is introduced in conversations when traffic is redirected. Ingress and egress traffic statistics, drops, congestion, and latency values are also represented in the traffic analytics statistics. This feature is supported on the Classic LAN, Enhanced Classic LAN, VXLAN, and BGP fabrics. For more information, see the "View traffic analytics for service endpoints" section in [View traffic analytics for endpoints](#).

Traffic analytics support for L3 external filter rules

By default, Nexus Dashboard enables traffic analytics on all L3Out interfaces to monitor traffic flows. You can create rules and apply them to external interfaces as an additional option to choose traffic that needs monitoring.

Applying interface rules can filter traffic collection and prevent overloading. It also allows for visibility into specific subnets and enables targeted monitoring of key traffic. For more information on applying interface filter rules, see [Apply filters to restrict traffic analytics for traffic analytics L3 external filtering](#).



- Data Center Interconnect (DCI) rules are not required for monitoring traffic within a fabric. They are only needed for monitoring traffic between multiple ACI sites. When you enable DCI rules for the first time, the system automatically creates a default rule on the spine Inter-Pod Network and Inter-Site Network (IPN/ISN) interfaces with the subnet 0.0.0.0/00::0/0. You can update these subnets to monitor only the traffic that needs monitoring.

- Egress TCAM carving must be configured on NX-OS fabrics when using the interface filter rule feature. For detailed instructions, refer to the "Performing Egress NetFlow TCAM Carving" section in the [Cisco N9000 Series NX-OS System Management Configuration Guide](#).

This configuration is required the first time interface filters are created. If Egress TCAM carving is not performed, the interface filter rule configuration fails on the switch, and the issue is reported as a configuration failure anomaly.

- Configuring interface filter rules will apply in the hardware/ASIC, this will also limit third-party NetFlow monitoring on that interface.

Guidelines and limitations for traffic analytics

- [General guidelines for traffic analytics](#)
- [Support information for traffic analytics with any fabric type](#)
- [Guidelines and limitations for traffic analytics compatibility mode](#)
- [Support information for traffic analytics with Cisco ACI fabrics](#)
- [Limitations for traffic analytics with Cisco ACI fabrics](#)
- [Support information for traffic analytics with NX-OS fabrics](#)
- [Limitations for traffic analytics with NX-OS fabrics](#)
- [Limitations for traffic analytics with multi-fabric and external conversations](#)
- [Limitations for traffic analytics with fabric groups](#)

General guidelines for traffic analytics

These guidelines apply for traffic analytics:

- Ensure that you have configured NTP, enabled PTP on the fabric, and PTP is synchronized within the fabric and between the fabric and the cluster. See the [Cisco Nexus Dashboard Deployment and Upgrade Guide](#) and [Precision Time Protocol \(PTP\) for Cisco Nexus Dashboard Insights](#) for more information.

Support information for traffic analytics with any fabric type

- For traffic analytics hardware support, refer to the **Supported fabric hardware/Supported 9000-series switches** area in the [Nexus Dashboard Capacity Planning tool](#).
- Traffic analytics does not support multicast.
- IPv6 underlay is not supported.
- L4-L7 traffic details are only supported and displayed for Layer 3 traffic. These details are not supported for Layer 2 bridged traffic.
- For traffic flowing from ACI to an NX-OS external network, the physical interface name might not be available when the traffic traverses the ACI border gateway tunnel. This behavior is due to a switch platform limitation in the ACI fabric export path.

Guidelines and limitations for traffic analytics compatibility mode

These guidelines and limitations apply for traffic analytics compatibility mode:

- NetFlow configuration requirement for Traffic Analytics compatibility mode on ACI fabrics

When enabling Traffic Analytics (TA) compatibility mode for ACI fabrics, you must configure the NetFlow configuration on the fabric with five tuple parameters - source IP, source port, destination IP, destination port, and protocol.

When the IP Protocol parameter is not configured, the output of the show flow cache command on switch devices will display the Protocol ID as 0. In this scenario, Nexus Dashboard Traffic Analytics Compatibility mode will not be able to classify the flow and there will be no data displayed under **Analysis Hub → Traffic Analytics**.

- In a multi-fabric configuration, you can configure each fabric for either **Traffic analytics full** mode or **Traffic analytics compatibility** mode. However, within a single fabric, mixed mode (where some switches operate in full mode and others in compatibility mode) is not supported.
- The following features are not supported in **Traffic analytics compatibility** mode:
 - Flow troubleshoot workflows
 - Traffic analytics support for traffic redirection to L4-L7 services
 - Traffic analytics support for L3 external filter rules
 - Latency and drop anomalies are not supported
- In **Traffic analytics compatibility** mode, for NX-OS fabrics, the recommended template data timeout value is 300 seconds, and the flow records timeout value is 60 seconds.

We recommend configuring the flow export timeout to 60 seconds to view the congestion anomaly.

- If **Traffic analytics compatibility** mode is disabled before the upgrade and then enabled afterward, the **Nexus Dashboard > Traffic Analytics** page will not display any conversations. To view the conversations, disable **Telemetry** at the fabric level, and then re-enable both **Telemetry** and **Traffic analytics compatibility** mode, re-add any custom categories in the **Manage service categories**.
- To switch from **Traffic analytics compatibility** mode to full mode, you must upgrade the controller or switches to the supported versions. **Traffic analytics full** mode requires ACI release 6.1.1f or later, or NX-OS version 10.4 or later. **Traffic analytics compatibility** mode supports versions earlier than ACI 6.1.1f and NX-OS 10.4.
- The following table lists the number of sites supported for Traffic analytics compatibility mode

Form factor	Recommended site limit
1 virtual Nexus Dashboard (vND)	1 site with default and additional custom categories.
1 physical Nexus Dashboard (pND)	2 sites with defaults and additional custom categories.

Form factor	Recommended site limit
3 data virtual Nexus Dashboard (vND)	1 site with default and additional custom categories.
3 physical Nexus Dashboard (pND)	3 sites with defaults and additional custom categories.
6 physical Nexus Dashboard (pND)	14 sites with defaults and additional custom categories.

Support information for traffic analytics with Cisco ACI fabrics

This support information applies for Traffic Analytics with Cisco ACI fabrics:

- Traffic analytics is supported on Cisco Application Policy Infrastructure Controller release 6.1(1f) and later release. Fabrics running Cisco Application Policy Infrastructure Controller version earlier than release 6.1(1f) can enable compatibility mode.
- Traffic analytics support for traffic redirection to Layer 4 to Layers 7 (L4-L7) feature requires Cisco Application Policy Infrastructure Controller release 6.1(2g) or later releases.
- Traffic analytics support for UDP service endpoints requires Cisco Application Policy Infrastructure Controller release 6.1(2g) or later releases.
- Traffic analytics support for L3 external filter rules requires Cisco Application Policy Infrastructure Controller 6.1(4) release.

Limitations for traffic analytics with Cisco ACI fabrics

- Traffic analytics for PBR may display zero latency and doubled traffic values.
- Traffic analytics may have incomplete information when virtual firewalls and endpoint VMs share the same DVS uplinks within a VMM domain.
- Traffic analytics may display inaccurate flow path visualization when chained L4-L7 service devices (for example, firewalls) are attached to the same Bridge Domains (BDs) on both the consumer and provider interfaces.

Support information for traffic analytics with NX-OS fabrics

This support information applies for traffic analytics with NX-OS fabrics:

- Traffic analytics full mode is supported on Cisco NX-OS release 10.4(2)F and later. Fabrics running NX-OS version earlier than the Cisco NX-OS release 10.4(2) can enable Traffic analytics compatibility mode.
- Traffic analytics UDP service endpoints and north-south interface filter features requires Cisco NX-OS release 10.5.(2) or later.
- Traffic analytics configurations or export are not supported on the Cisco N9500 modular chassis, but flow troubleshoot jobs are supported for FX platform switches and Cisco N9500 modular chassis.

Limitations for traffic analytics with NX-OS fabrics

These limitations apply for traffic analytics with NX-OS fabrics:

- Traffic analytics may display partial data when the VRF instance is configured with the new L3VNI mode. For more information about the new L3VNI mode, see the [Cisco N9000 Series NX-OS VXLAN Configuration Guide](#).
- Traffic analytics configurations or export is not supported on Cisco N9500 modular chassis, however flow troubleshoot jobs is supported for FX platform switches and Cisco N9500 modular chassis.
- You can apply traffic analytics L3 external filters in VXLAN fabrics for VRF-Lite external interfaces only, and switches must have Cisco NX-OS release 10.5(3) or later.
- The timestamp tagging (TTAG-strip) should be configured on DCI interfaces. Nexus Dashboard displays only the inter-fabric flow latency, that is, the latency between the border router and the internal flow endpoint. The end-to-end latency is not displayed in this case.

Limitations for traffic analytics with multi-fabric and external conversations

These limitations apply for traffic analytics with a multi-fabric configuration:

- Flow troubleshoot jobs are local to a fabric.
- The flow troubleshoot workflow is not supported for transit conversations.
- Latency computation is local to the fabric. The computed latency is from the endpoint to the border node and not end-to-end.
- Remote endpoints with overlapping IP addresses are not enriched with remote fabric properties.
- Traffic analytics does not show conversations between external endpoints. One endpoint in the conversation should be learned or local to the fabric. This applies for both a single fabric and a multi-fabric.
- Metadata visibility in multisite cross-VRF flows: In ACI multisite environments, for traffic flows between ESG and L3Out, the service endpoint details do not display client metadata, such as the tenant, VRF instance, and EPG, if:
 - The client endpoint is external (behind an L3Out).
 - The client endpoint is on a remote site relative to the service.
 - The client and service endpoints are in different VRF instances.

In this scenario, these fields are empty or contain em dashes (–).

Limitations for traffic analytics with fabric groups

These limitations apply for traffic analytics with fabric groups:

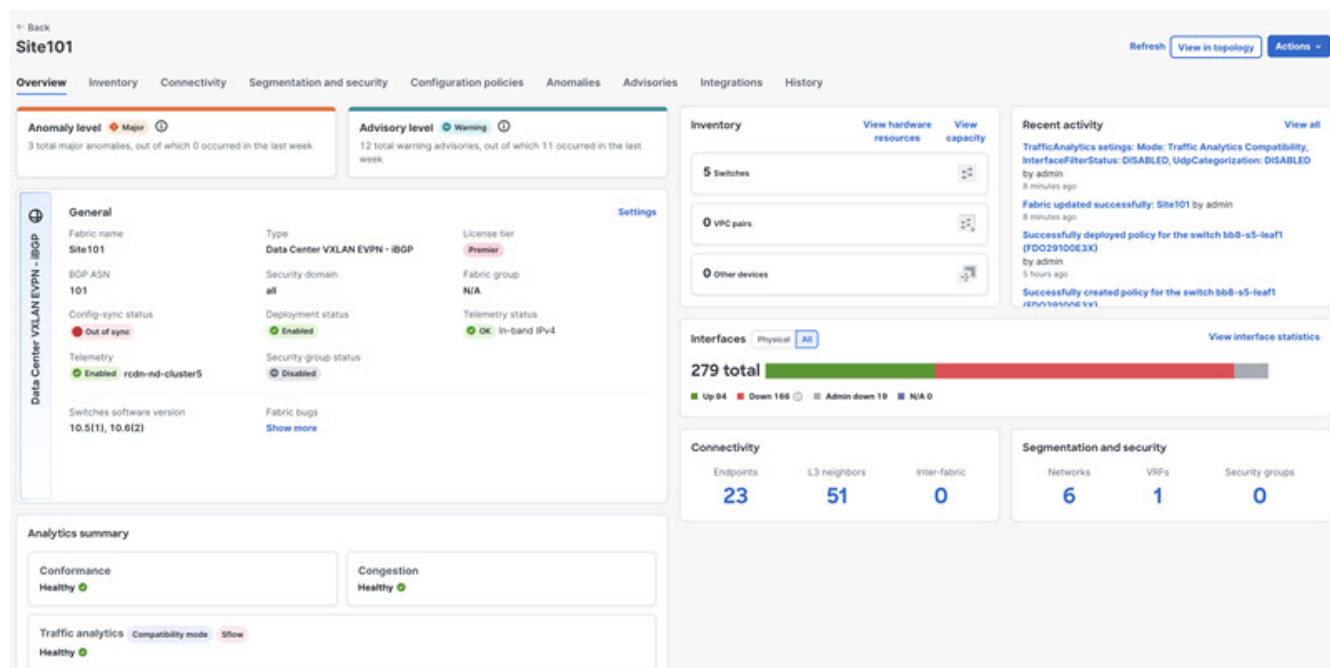
- When viewing traffic analytics for a fabric group, as described in [View traffic analytics for a fabric group](#), all member fabrics must have the same Category name, protocol and port mappings; otherwise the traffic analytics information displayed for the fabric group will have incorrect information.

Traffic analytics support for Silicon One (S1) platforms

Beginning with Cisco Nexus Dashboard 4.3.1, **Traffic Analytics Compatibility** mode is supported on Netflow-based Silicon One (S1) platforms, running NX-OS 10.6(3) release or later, in addition to sFlow-based N9K-C9804, N9K-C9808, and G200 switches.

Nexus Dashboard 4.3.1 provides a baseline level of Traffic analytics feature support on S1-based switches by enabling Traffic analytics compatibility mode with the appropriate flow export (Netflow or sFlow) method.

The following screenshot shows Traffic analytics compatibility mode sflow configured for the fabric.



Enhancements to Traffic analytics compatibility mode

In Nexus Dashboard 4.3.1, Traffic analytics compatibility mode includes the following enhancements for NX-OS fabrics:


- Traffic analytics compatibility mode support is extended to additional S1 platforms.
- Netflow-based Traffic analytics compatibility support is available for supported Krypton-based platforms.
- sFlow-based Traffic analytics compatibility support is introduced for supported N9K-C9804, N9K-C9808, and G200 platforms.
- When you configure **Traffic analytics**, Nexus Dashboard chooses one of these Traffic analytics compatibility submodes:
 - Netflow
 - sFlow
- The **System Status > Telemetry** page also reflects the active Traffic analytics modes for the fabric in the **TA mode** column. Depending on the traffic analytics configuration mode, you will see one of these values in the **TA mode** column:
- **Full**

- Compatibility mode netflow
- Compatibility mode sflow



Traffic analytics compatibility mode with sFlow is not supported on ACI fabrics.

The following screenshots display enhancements to the **Traffic Analytics > Overview** page. The Traffic Analytics Compatibility mode and the flags for the configured Sub-modes (Sflow or Netflow) appear on top of the **Overview** page.



Traffic Analytics Services

View Service categories by Traffic analytics score

Filter by attributes

View endpoint details

May 11, 2026, 11:25 AM

Endpoint	Hostname	Service port	Conversation attributes	Location	Tenant	VLAN	VRF	Switches	Interface	Fabric	Security group	Traffic analytics score	Congestion score
20.1.1.12	-	80	Local	Local	default	3	vr101	bb8-s5-leaf3	eth1/2/1	Site101	-	Healthy	Health
2020:12	-	80	Local	Local	default	3	vr101	bb8-s5-leaf3	eth1/2/1	Site101	-	Healthy	Health

Supported switches and behavior

The following switches are supported in 4.3.1:

- Netflow based switches:

- N9396Y12C-SE1
- N9324C-SE1U
- N9348Y2C6D-SE1U
- N9336C-SE1
- N9396T12C-SE1
- sFlow-based switches:
 - N9K-C9804, N9K-C9808
 - N9364E-SG2-Q, N9364E-SG2-O



- The **Congestion** anomaly support depends on the switch capability and whether the switch exports the required CWR flag.
- Byte count is not supported on the Netflow based switch models. The show flow cache stores up to 1,000,000 flows. However, due to ASIC limitations, the byte count displayed in the output of this command is always zero (0), and the collector also receives a byte count of zero (0). See the [Nexus Configuration Guide](#) for more information.
- If an NX-OS fabric contains a mix of CS and S1 (E100 or G200) switches, PTP is not supported and flow telemetry cannot be enabled.

Traffic analytics compatibility mode selection

Switch combinations	Supported Traffic analytics mode	Behavior
N9K-C9804 and N9K-C9808 switches on spine + (CS+E100) on leaf	Traffic analytics compatibility with Netflow	E100 dictates Netflow; CS adapts; N9K-C9804 and N9K-C9808 switches are ignored
N9K-C9804 and N9K-C9808 switches on spine + (CS+G200) on leaf	Traffic analytics compatibility with sFlow	G200 dictates sFlow, CS adapts, N9K-C9804 and N9K-C9808 switches are ignored
CS + E100 + G200 on leaf	Not supported	E100 requires Netflow, G200 requires sFlow - incompatible
E100 + G200 on leaf	Not supported	E100 requires Netflow, G200 requires sFlow - incompatible
(E100 + N9K-C9804 and N9K-C9808 switches) on leaf	Traffic analytics compatibility with Netflow	N9K-C9804 and N9K-C9808 switches as leaf are not validated
(G200 + N9K-C9804 and N9K-C9808 switches) on leaf	Traffic analytics compatibility with sFlow	N9K-C9804 and N9K-C9808 switches as leaf are not validated
(CS + N9K-C9804 and N9K-C9808 switches) on leaf	Traffic analytics full (primary) + both compatibility mode protocols available	N9K-C9804 and N9K-C9808 switches as leaf are not validated

Supported topologies

In addition to the existing Cloud Scale (CS) switch topology, where all switches are CS switches, Nexus Dashboard 4.3.1 now supports these additional topologies:

S1 switches and supported Traffic analytics modes

Topology: Leaf or Host facing	Enabled Feature from switch side	4.3.1 support	Traffic analytics mode
All E100 switches	Netflow	Netflow supported	Traffic analytics compatibility
All G200 switches	sFlow/sampled NF	sFlow supported	Traffic analytics compatibility
CloudScale + E100 switches	Netflow	Netflow supported	Traffic analytics compatibility
CloudScale + G200	sFlow	sFlow supported	Traffic analytics compatibility
CloudScale + E100 + G200	sFlow	Not supported	None
N9K-C9804 and N9K-C9808 switches (cannot be host facing) + CS/E100/G200	sFlow	Not supported	None

Upgrade and migration

If you upgrade from Nexus Dashboard 4.1.1g or 4.2 to 4.3.1, and the fabric was previously running in Traffic analytics compatibility mode, you should verify the configured compatibility mode after the upgrade.



Earlier releases defaulted Traffic analytics compatibility to Netflow. In 4.3.1, because Traffic analytics compatibility mode can now run in either Netflow or sFlow mode, the correct mode must be explicitly selected to match the switch platform and deployment.

Guidelines and limitations for Traffic analytics support for S1 platforms

- Only **Congestion** anomalies are supported in **Traffic analytics compatibility** mode for known service endpoints.
- In environments with multiple fabrics, you can configure each fabric independently, but each fabric must remain in a single mode.
- When switching from **Traffic analytics full** mode to **Traffic analytics compatibility** mode:
 - You must push the required Netflow or sFlow configuration to the switches.
 - You must remove the compatibility mode flow export configuration from the switches.
- Mixed mode within a single fabric is not supported. In systems with multiple fabrics, each fabric

can independently run in Traffic analytics full or compatibility mode. But within a fabric, there is no support for mixed mode (some switches running in Traffic Analytics full mode and some in compatibility mode or some switches exporting Netflow and some exporting sFlow)

- Support details for some platform and software combinations may depend on the switch type and software version.
- Some anomaly and telemetry capabilities may vary by platform.
- 1-node virtual Nexus Dashboard (Data node) is not recommended for AI fabrics with sFlow enabled. The processing capacity of a single-node virtual deployment is insufficient to handle the sFlow sample rate generated by high-speed AI fabric workloads. If you observe sFlow record drops, increase the sFlow sampling rate to 1:100,000. A higher sampling rate reduces the number of samples generated per link.

Configure traffic analytics flow collection

Before you begin

- Ensure traffic analytics is enabled for your cluster in the **Admin > System Settings > Flow collection**.

To enable traffic analytics, choose **Traffic analytics** in the **Flow collection mode**.

- Review the guidelines and limitations described in [Guidelines and limitations for traffic analytics](#) for the traffic analytics mode you choose for your fabric.
- If you are configuring **Traffic analytics compatibility** mode, review the support information in [Guidelines and limitations for traffic analytics compatibility mode](#) section.



If flow telemetry is already enabled on the fabric, you must first disable flow telemetry for all the fabrics and remove all flow rules before enabling traffic analytics.

Apply the traffic analytics configuration

Follow these steps to apply the traffic analytics configuration.

1. Navigate to the main **Fabrics** page.

Manage > Fabrics

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

The **Edit *fabric_name* Settings** page appears.

4. Click **Telemetry > Configuration**.
5. Navigate to **Flow collections modes** under **Traffic analytics**.
6. Choose one of these modes.
 - **Traffic analytics full** to configure all traffic analytics features.

- o **Traffic analytics compatibility** to configure compatibility mode and follow the additional steps described in [Apply the traffic analytics compatibility mode configuration](#) section. Nexus Dashboard automatically chooses one of these **Sub-modes** based on the switches present in your fabric:

- **NetFlow**
- **sFlow**

- o **Disabled** to turn off the traffic analytics configuration on your fabric.

7. Click the **Traffic analytics scope** drop-down list to choose one of these scopes.

- o **Inter-fabric + external** to monitor all traffic regardless of whether the client and/or service endpoint is learned in the local fabric or externally.
- o **Inter-fabric** to monitor all traffic as long as the either the client or the service endpoints are learned in the local fabric.
- o **Intra-fabric** to monitor traffic where both client and service endpoints are learned in the local fabric.

8. Click **Save** to save your changes.

Navigate to your fabric's **Overview** page to review the traffic analytics mode information for your fabric. Depending on the traffic analytics mode you configured, the traffic analytics status (enabled or disabled) and summary are available in the **Analytics summary > Traffic analytics**. For more information, see [View traffic analytics](#).

You can view the traffic analytics data for your fabrics after applying the traffic analytics configuration.

- o In the **Conversation rate status** area, you can see the number of conversations that exceed the limit and traffic analytics drops.

You must make sure that you do not exceed the maximum conversation rate limit. If you exceed the maximum conversation rate limit, you will see drops in flows records, which will impact the visibility.

- o Click **View all traffic analytics rate statistics** to view the statistics for each switch in a fabric.

Click **Enabled** for this feature if you want to track multi-fabric and external conversations (north to south).

Apply the traffic analytics compatibility mode configuration

You must manually configure NetFlow to apply the traffic analytics compatibility mode configuration on your NX-OS and Cisco ACI fabrics.

Follow these steps to apply the traffic analytics compatibility mode configuration.

For NX-OS Fabrics:

1. Navigate to **Admin > System Status > Telemetry**.
2. Click the **Switches** tab in the **Telemetry Status** area.
3. Click the ellipsis icon for the NX-OS 9000 switch and then click **Expected Configuration**.

The **Expected Configuration for *switch_name*** page appears.

4. View and copy configurations under **Software Telemetry** and **Flow Telemetry**.
5. Using the command line, log in to the switch.
6. Paste the configuration copied in step 4 to the switches.



The configuration shown in the **Expected Configuration for *switch_name*** is a sample configuration. You must customize this configuration before applying it to the switch. For more information, see [NetFlow](#) documentation.

For Cisco ACI fabrics:

- Follow the configuration details described in [Cisco APIC and NetFlow](#).
- Configure NetFlow from the APIC with the exporter IP addresses as the Nexus Dashboard's collector persistent IP addresses. Navigate to **Admin > System Settings > External pools** to see the Nexus Dashboard's collector persistent IP addresses. The port number should be **5641**.
- Verify that you have the NetFlow priority configured under the fabric node control policy in APIC.

View flow collection status

For NX-OS fabrics, you can view the **Flow Collection** status for each node in the **Switch Configuration Status** column.

- Green—Flow collection is successfully enabled.
- Red—Flow collection is not enabled.
- Orange—Flow collection is partially enabled.
- Gray—Flow collection is not supported or data cannot be found. If a switch is in the disabled state, it is included in the gray category.

View traffic analytics

You can view traffic analytics at these levels.

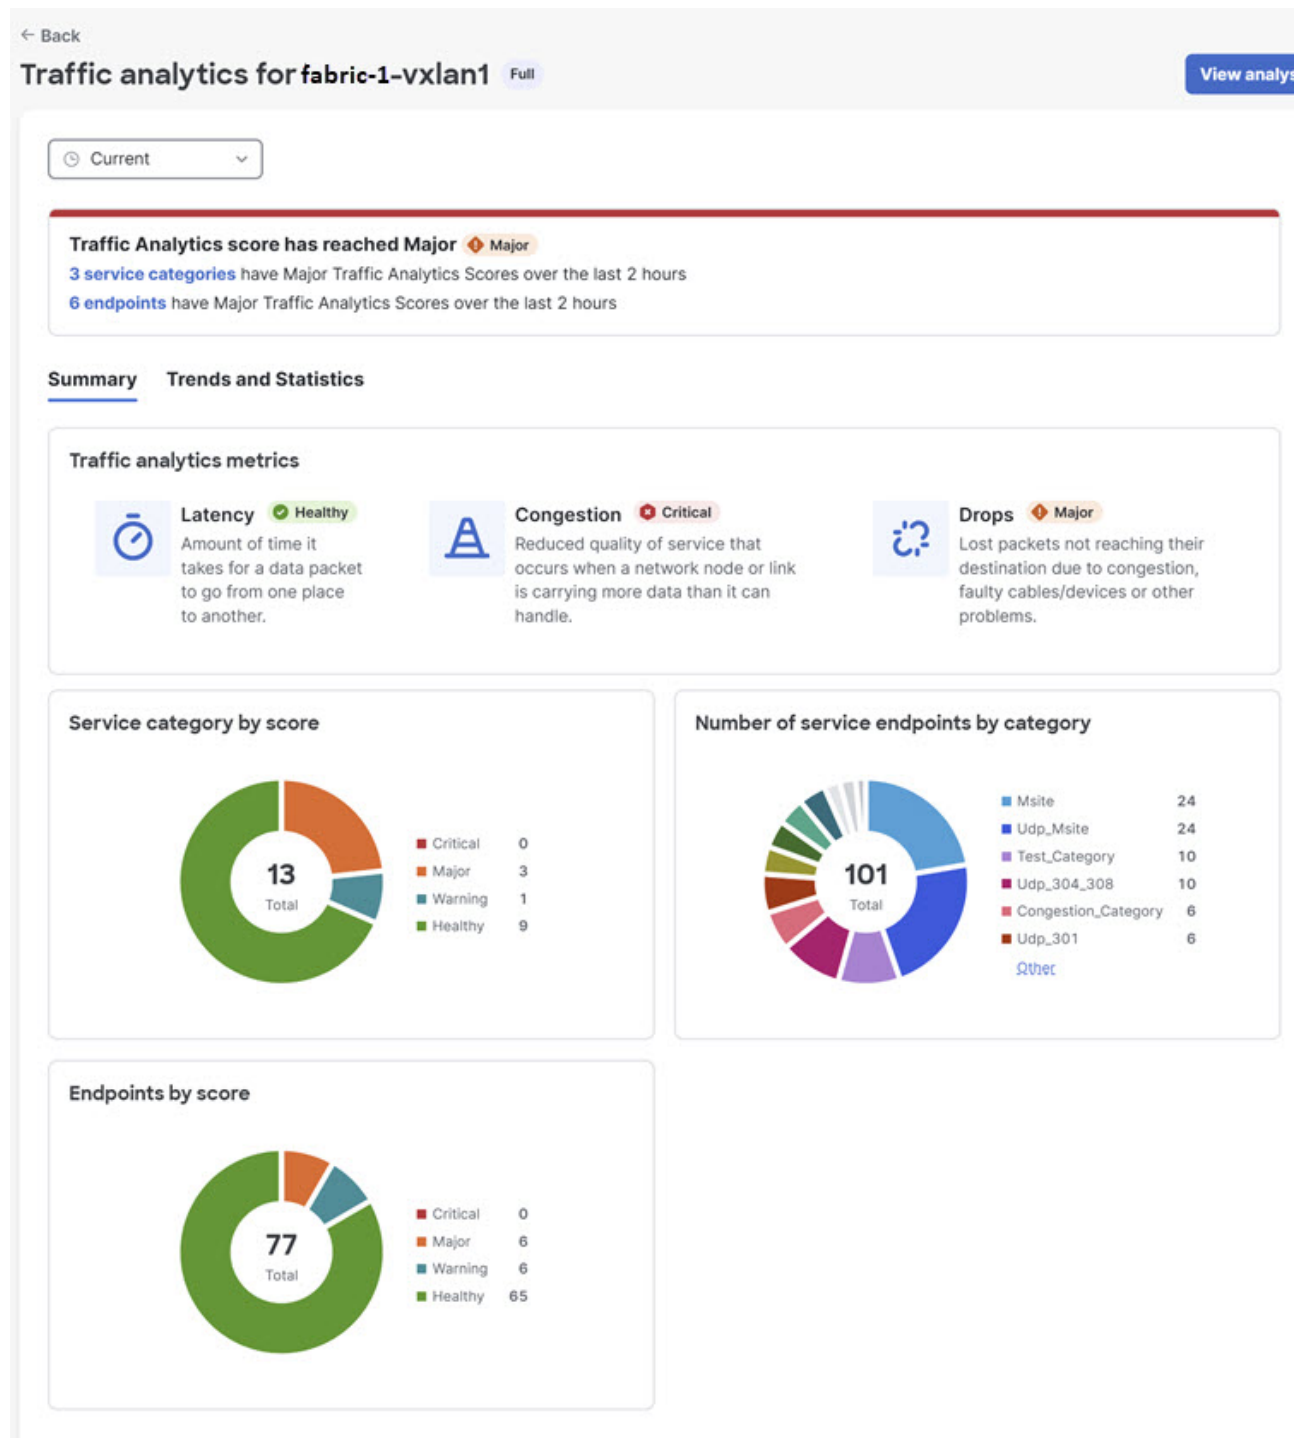
- [View traffic analytics for an individual fabric](#)
- [View traffic analytics for a fabric group](#)
- [View traffic analytics status for all fabrics](#)

View traffic analytics for an individual fabric

Follow these steps to view Traffic analytics for an individual fabric.

1. Navigate to **Manage > Fabrics** page.
2. Click the fabric name.
3. In the **Analytics summary** area, click **Traffic analytics** to view traffic analytics details for that fabric.

In the **Traffic Analytics** page, all information is grouped as service categories for that fabric.



4. Choose a time range from the drop-down list.

By default the last 2 hours from the current time is chosen.

5. The **Summary** area displays the **Traffic analytics metrics** and how the metrics is determined.

6. Click **Trends and Statistics** to view **Traffic profile**, **Top endpoint service score changes**, and **Top endpoint categories by**.

🕒 Current


Traffic Analytics score has reached Major 🔴 Major

3 service categories have Major Traffic Analytics Scores over the last 2 hours

6 endpoints have Major Traffic Analytics Scores over the last 2 hours

Summary

Trends and Statistics



Top endpoint service score changes
over the last 2 hours

Categories	Score change	Highest severity	Affecting metric
Congestion_Category	🟡 Warning → 🟡 Warning	🟡 Warning	Congestion →
Drop_Category	🔴 Major → 🔴 Major	🔴 Major	Drop →
Latency_Category	🟢 Healthy → 🟢 Healthy	🟢 Healthy	-
Login_Uilities	🟢 Healthy → 🟢 Healthy	🟢 Healthy	-
Msite	🟢 Healthy → 🟢 Healthy	🟢 Healthy	-
Test_Category	🟢 Healthy → 🟢 Healthy	🟢 Healthy	-
Udp_301	🟢 Healthy → 🟢 Healthy	🟢 Healthy	-
Udp_304_308	🟢 Healthy → 🟢 Healthy	🟢 Healthy	-
Udp_311	🟢 Healthy → 🟢 Healthy	🟢 Healthy	-
Udp_Drop	🔴 Major → 🔴 Major	🔴 Major	Drop →

Top endpoint categories by Rx latency 📄
over the last 2 hours (Averages and Trends)

Categories	Highest Value	Trend
Latency_Category	3 μs	→
Login_Uilities	3 μs	→
Test_Category	3 μs	→
Udp_301	3 μs	→
Udp_304_308	3 μs	→
Udp_311	3 μs	→
Web	3 μs	→
Congestion_Category	0 μs	→
Drop_Category	0 μs	→
Msite	0 μs	→

- a. In the **Traffic profile** area you can view the traffic amount for the endpoint service category.
- b. In the **Top endpoint service score changes** area, you can view the anomaly score change across the chosen time range and the metrics (such as latency, congestion, or drops) affecting the score change.
- c. In the **Top endpoint categories by area** you can see the top categories by **Rx and Tx Latency, Congestion, and Drop**.

7. Click **View analysis** to view traffic analytics for all the fabrics.


View traffic analytics for a fabric group

You can review traffic analytics for a fabric group in the Nexus Dashboard, which displays aggregated information from the member fabrics in the fabric group.

Follow these steps to view traffic analytics status for a fabric group.

1. Navigate to **Manage > Fabrics > Fabric groups** page.
2. Click the fabric group name.
3. In the **Analytics summary** area, click **Traffic analytics** to view traffic analytics details for that fabric group.

In the **Traffic Analytics** page, all information is grouped as service categories for that fabric group.



4. Choose a time range from the drop-down list.

By default the last 2 hours from the current time is chosen.

5. The **Summary** area displays the **Traffic analytics metrics** and how the metrics is determined.
6. Click **View analysis** to view traffic analytics for all the fabrics.
7. Click one of these tabs to view traffic analytics for the fabric group at these levels.

- [Overview](#)
- [Services](#)
- [Endpoints](#)

You can also view traffic analytics for a fabric group by navigating through the **Analysis Hub** page.

1. Navigate to **Analyze > Analysis Hub > Traffic analytics**.
2. Click the first drop-down list and choose the fabric group that you want to run traffic analytics on.

The **Traffic Analytics** page is displayed for that fabric group.

3. Choose a time range from the second drop-down list.

By default the last 2 hours from the current time is chosen.

4. Click one of these tabs to view traffic analytics for the fabric group at these levels.

- [Overview](#)
- [Services](#)
- [Endpoints](#)

Overview

The **Overview** area displays the traffic analytics score and how the metrics are determined in these tiles:

- **Summary**
- **Traffic analytics metrics**
- **Service category by score**

Services

The **Services** page displays the service endpoints across all member fabrics in the selected fabric group, where the member fabric is shown in the **Reporting fabric** column in the table, while the **Fabric** column indicates where the endpoint has been learned.

If the same service endpoint is reported by multiple fabrics, then it is displayed as multiple entries.

Service endpoint details

For a specific service:

- The **Overview** traffic analytics anomaly score is the maximum across all fabrics in the fabric group.
- **Anomalies**: Shows anomalies for all fabric group member fabrics for this service.
- **Flows**: Shows flow troubleshoot jobs that are run in the member fabrics in the fabric group for this service endpoint.

Endpoints

The **Endpoints** page displays the endpoints across all member fabrics in the selected fabric group, where the member fabric is shown in the **Reporting fabric** column in the table, while the **Fabric** column indicates where the endpoint has been learned.

If the same endpoint is reported by multiple fabrics, then it is displayed as multiple entries.


Endpoint details

For a specific endpoint:

- The **Overview** traffic analytics score is the maximum across all fabrics in the fabric group.
- Connections to other services and IP addresses from this endpoint: Conversation data is merged across fabrics.
- Anomalies: Endpoint anomalies for all member fabrics are displayed. The fabric shown in the **Fabric** column identifies where the anomaly originated.
- Flows: Troubleshoot jobs for all member fabrics are displayed.

View traffic analytics status for all fabrics

You can review traffic analytics for all fabrics in the Nexus Dashboard **Admin > System Status** section.



The screenshot shows the 'System Status' page with the 'Telemetry' tab selected. Under 'Telemetry status', the 'Fabrics' sub-tab is active. A 'Traffic analytics' section contains a table with the following data:

Fabrics	TA mode	TA capability	Telemetry config status	Software telemetry	Flow collection	Software telemetry switch status	Flow col
ifav22-site1	Full - Inter-fabric + external	11 0	OK	Enabled	Enabled	-	-
ifav22-site2	Full - Inter-fabric + external	2 0	Not OK	Add failed	Enable failed	-	-

Follow these steps to view traffic analytics status for all fabrics.

1. Navigate to **Admin > System Status > Telemetry**.
2. Click the **Fabrics** tab to view the telemetry details for all fabrics. For more information, see the Telemetry section in [Reviewing System Status for Your Nexus Dashboard](#).

Manage service categories

In the **Manage Service Categories** area, you can view the ports that have been assigned to categories based on standard networking defaults and any categories you may have created. If a port has not been assigned to a category, you can assign it to one of the existing categories or create a new category. This helps you to organize and manage your network ports more efficiently.

You can view the UDP conversations in the traffic analytics table along with the TCP conversations. You can create a new category or modify an existing category and choose UDP or TCP protocol to view and monitor the conversations.

< Manage Service Categories

×

Below are ports that have been assigned to categories based on standard networking defaults and any categories you may have created. If a port has not been assigned to a category, you can assign it to one of the existing categories or create a new category. This helps you to organize and manage your network ports more efficiently

Known ports on categories

Filter by attributes		New category
Category name	Port selectors	
Analytics	TCP 3000 TCP 4739-4740 TCP 5000 TCP 5601	...
	UDP 4739-4740	
Database	TCP 66 TCP 118 TCP 591 TCP 1186 TCP 1433-1434	...
	TCP 2399 TCP 2483-2484 TCP 3050 TCP 3306 TCP 5432	
Distributed_DataStore	TCP 5984 TCP 7474 TCP 8000 TCP 8091 TCP 8999-9000	...
	TCP 9042 TCP 9090 TCP 9200 TCP 9300 TCP 27017	
Email	UDP 66 UDP 118 UDP 591 UDP 1186 UDP 1433-1434	...
	UDP 2399 UDP 2483-2484 UDP 3050 UDP 3306 UDP 5432	
File_Sharing	UDP 5984 UDP 8000 UDP 8091	...
	TCP 76 UDP 76	
Gaming	TCP 24-25 TCP 50 TCP 58 TCP 102 TCP 109-110 TCP 143	...
	TCP 220 TCP 465 TCP 587 TCP 691 TCP 993 TCP 995	
	UDP 24-25 UDP 50 UDP 58 UDP 102 UDP 109-110 UDP 220	...
	UDP 587 UDP 691 UDP 995	
	TCP 20-21 TCP 59 TCP 69 TCP 115 TCP 119 TCP 445	...
	TCP 563 TCP 873 TCP 989-990 TCP 1214 TCP 2049	
	TCP 5402 UDP 20-21 UDP 59 UDP 69 UDP 115 UDP 119	...
	UDP 445 UDP 563 UDP 873 UDP 989-990 UDP 1214	
	UDP 2049 UDP 5402	...
	TCP 1725 TCP 3071 TCP 3074 TCP 3724 UDP 1725	...
	UDP 3071 UDP 3074 UDP 3724	

Follow these steps to manage service categories.

1. Navigate to **Analyze > Analyze Hub > Traffic analytics**.
2. Choose a fabric from the drop-down list.
3. In the **Service category by score** area, click **Manage > Fabrics > Edit Fabric Settings**.
4. Click **Manage service categories**.

The **Manage Service Categories** page appears.

5. Click **Actions > New category** to create a new category.

Manage service categories

New service category
Category name *

Port selectors

Protocol	Ports	
TCP	70,80,443,8080	
UDP	70,80,443,8080	

+ Add

Cancel

Save

6. Enter the name of the category.
7. From the **Protocol** drop-down list, choose **TCP** or **UDP** protocol.
8. In the **Ports** field, enter a port number or range.
9. Click **Add** to add additional protocols.
10. Click **Save**.
11. To edit a category, click the ellipsis icon and choose **Actions > Edit**.
 - a. Edit the values and click **Save**.
12. To delete a category, click the ellipsis icon and choose **Actions > Delete**.
 - a. Click **Confirm**.

Apply filters to restrict traffic analytics for traffic analytics L3 external filtering

Follow these steps to apply filters to restrict traffic analytics for traffic analytics L3 external filtering.

1. Navigate to the main **Fabrics** page.

Manage > Fabrics

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

The **Edit *fabric_name* Settings** page appears.

4. Click **Telemetry > Configuration**.
5. Navigate to **Flow Collections Modes** under **Traffic analytics**.
6. Choose **Traffic analytics full** mode.
7. Check the **Traffic analytics filtering rules** check box.

Edit Fabric-site2 Settings

i

Some fabric settings can only be modified from the ND Cluster owner Fabric-cluster

General
Telemetry
External streaming

Configuration
NAS

General

☒

Enable assurance analysis ⓘ

☐

Enable microburst sensitivity

Microburst sensitivity level*

Medium sensitivity
▼

Traffic analytics

Flow collection modes ⓘ

☒ Traffic analytics full ⓘ

☐ Traffic analytics compatibility ⓘ

☐ Disabled

Traffic analytics scope ⓘ

Inter-fabric + external
▼

☒

Enable UDP categorization

Service categories

Manage service categories

☒

Traffic analytics filtering rules ⓘ

i

System generated rules are created for all detected L3 connections permitting all subnets (0.0.0.0/0). To restrict Traffic analytics flows, modify existing rule with a subnet prefix, or create a new rule.

Filter by attributes
Actions ▼

☒	Rule	Rule status	Tenant	L3Out	Switch	Interfaces	DCI/ISN	Subnets	
☒	DCI	⏸ Disabled	infra	intersite	Fabric -spine7	eth1/16, eth1/14, eth1/10, eth1/13, eth1/15	True	0.0.0.0/0, ::/0	

8. Check the check box next to the filter **Rule**. If you do not have an existing filter rule, you can create a rule, see [Create an interface filter rule](#) for details.

9. Click **Save**.

Create an interface filter rule

Follow these steps to create an interface filter rule.

1. Navigate to the main **Fabrics** page.

Manage > Fabrics

2. Locate the fabric you want to edit.

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

The **Edit *fabric_name* Settings** page appears.

4. Click **Telemetry > Configuration**.
5. Navigate to **Flow Collections Modes** under **Traffic analytics**.
6. Choose **Traffic analytics full** mode.
7. Check the **Traffic analytics filters rules** check box.
8. Click the **Actions** drop-down list and choose **Create**.

The **Create traffic analytics filter rule** page appears.

9. Enter the appropriate value for the fields.



The field information for creating filters differs depending on the fabric for which you are creating filter rules.

- o For ACI fabrics, enter a value for these fields.
 - **Rule Name**—Specify a name for the filter rule.
 - **DCI/ISN**—Check the check box if this is a Data Center Interconnect (DCI) and Intersite Network (ISN) connectivity. This field is only applicable for ACI fabrics.
 - **Enable Rule**—Check the check box to enable this rule for your fabric.
 - **Tenant**—Click on the plus sign to add these tenant field values.
 - **L3Out**—Choose the Layer 3 out or external-routed network.
 - **Switch**—Specify the switch name.
 - **Interfaces/VLAN**—Specify the interfaces/VLAN information.
 - **Subnet**—Enter one or more subnet IP addresses and press the **Enter** key from your keyboard.
- o For NX-OS fabrics, enter a value for these field.



You must create VRF Lite connectivity before creating the filter rules. For more information on VRF Lite connectivity, see the "Establishing inter-fabric connectivity using VRF Lite" section in [Working with Connectivity in Your Nexus Dashboard LAN Fabrics](#).

- **Rule Name**—Specify a name for the filter rule.
- **Enable Rule**—Check the check box to enable this rule for your fabric.
- **VRF**—Click the plus sign and choose a VRF name from the drop-down list.
- **Switch**—Choose a switch name from the drop-down list.
- **Interfaces/VLAN**—Choose interfaces from the drop-down list, check the check box for each interface you want to add for this rule.
- Click on the check icon to save the information.
- Click the plus sign to add interfaces.

- **Subnet**—Enter one or more subnet IP addresses and press the **Enter** key from your keyboard.

10. Click **Save** to save this rule.

In addition to creating filter rules, you can perform edit (modify), disable, enable, or delete actions using the **Actions** drop-down list.

- o Check the check box next to the filter rule, click **Actions**.
 - Choose **Edit** to modify the filter rule. Note that you cannot modify the name of an existing rule.
 - Choose **Enable** to enable the filter rule or re-enable the rule that you disabled.
 - Choose **Disable** to disable the filter rule. Once the rule is disabled, the filter rule remains in the system but stays inactive.
 - Choose **Delete** to delete the filter rule for the fabric.

View traffic analytics for endpoints

Follow these steps to view Traffic analytics for endpoints.

1. Navigate to **Manage > Fabrics**.
2. Click the *fabric* name.
3. Navigate to **Connectivity > Endpoints**.
 - a. In the Endpoint table click an *IP address*.
 - b. In the **IP details** page, click **Traffic analytics** tab to display the Traffic analytics view for endpoints.

Overview Endpoint history Anomalies **Traffic analytics** Trends and statistics Flow collections

Traffic analytics score reached **Healthy** Healthy

This score change generated 0 anomalies over the last 2 hours [View anomalies](#)

Services hosted on this endpoint

Filter by attributes

Service port	Conversation attributes ⓘ	Location	VLAN	Traffic analytics score ▼	Category	Protocol	TCP SYN Count (Intra-fabric)	TCP RST Count (Intra-fabric)
8018	External, Local	Local	-	Healthy	Msite	TCP	6892	6892
8038	External, Local	Local	-	Healthy	Udp_Msite	UDP	0	0

Connections to other services and IPs from this endpoint by **Traffic analytics score** ▼ over the last 2 hours

Filter by attributes

[Start Flow collection](#)

Connections to other services and IPs from this endpoint by **Traffic analytics score** ▼ over the last 2 hours

Filter by attributes

[Start Flow collection](#)

Endpoint	Conversation attributes ⓘ	Service port	Interface	Traffic analytics score ▼	Category	Protocol	Tenant
☐ 20.11.0.21	External	8031	eth1/7/1	Healthy	Udp_Msite	UDP	-
☐ 20.11.0.21	External	8011	eth1/7/1	Healthy	Msite	TCP	-
☐ 20.13.0.23	Local	8033	eth1/54/1	Healthy	Udp_Msite	UDP	-
☐ 20.13.0.23	Local	8013	eth1/54/1	Healthy	Msite	TCP	-

View traffic analytics for service endpoints

You can view the service endpoint details from the **Analysis hub** page.

Follow these steps to view Service Endpoint details.

1. Navigate to **Analyze > Analysis hub > Traffic analytics**.
2. Click on the **Endpoints** tab.
3. Click on an IP address listed in the **Endpoint** column to display the Endpoint Details page.
4. Scroll down to the **Services hosted on this endpoint** area.
5. Click on a port number listed in the **Service port** column to view the **Service Endpoints Details** page.
6. In the **Service Endpoints Details** page, scroll down to view the **Endpoint, Conversation attributes, Location, L4-L7 attach, L4-L7 cluster**, and other field information.

The **Conversation attributes** column displays various flags; an "L4-L7" flag indicates the flow is passing through an L4-L7 device.

The **L4-L7 cluster** column displays the name of the device and interface details if the interface details are available. The **L4-L7 attach** column displays where it is attached.

Find traffic analytics conversations for an endpoint in a fabric

Traffic analytics conversations, traffic analytics records, and service endpoints can also be found from the fabric page by navigating to the endpoint details, which filters all the traffic analytics records that we have for that specific endpoint.

1. Navigate to the **Manage > Fabrics** page.
2. Click the *fabric* name.
3. In the fabric **Overview** page, navigate to **Connectivity > Endpoints**.

You can also navigate to the **Connectivity** area in your fabric's **Overview** page and click on **Endpoints** to launch the **Endpoints** page.

🕒 Current ▾

Summary

Anomaly level



■ Critical 0
 ■ Major 0
 ■ Minor 0
 ■ Warning 0
 ■ Healthy 15

Change summary ⓘ

0 Endpoint New or modified



0 Endpoint

Status == Active ✕

MAC	IP addresses	Hostname	Anomaly level	Switches	Interface	Status
00:11:22:00:58:02	20.18.0.28 20:18::28	-	✓ Healthy	ifav22-leaf14	eth1/54/3	✓ Active
00:11:22:00:53:02	20.13.0.23 20:13::211:22ff:fe00: - +1 More	-	✓ Healthy	ifav22-leaf14	eth1/54/1	✓ Active
00:11:22:00:54:02	20.14.0.24 20:14::211:22ff:fe00: - +1 More	-	✓ Healthy	ifav22-leaf14	eth1/54/1	✓ Active
04:BD:97:20:88:EF	10.6.0.14	-	✓ Healthy	ifav22-spine9	eth1/50.4	✓ Active
04:BD:97:20:88:EF	10.6.0.10	-	✓ Healthy	ifav22-spine9	eth1/49.20	✓ Active
2C:5A:0F:5D:B8:97	10.7.0.17	-	✓ Healthy	ifav22-leaf14	eth1/51.10	✓ Active
00:22:BD:F8:19:FF	20.2.208.1	-	✓ Healthy	ifav22-leaf14	eth1/53.4	✓ Active
F8:39:18:4D:41:BF	10.6.0.6	-	✓ Healthy	ifav22-spine9	eth1/52.4	✓ Active
F8:39:18:4D:41:BF	10.7.0.18	-	✓ Healthy	ifav22-spine9	eth1/52.10	✓ Active
2C:5A:0F:5D:B8:97	10.6.0.9	-	✓ Healthy	ifav22-leaf13	eth1/51.20	✓ Active
B0:8B:D0:81:6D:A7	10.6.0.17	-	✓ Healthy	ifav22-leaf13	eth1/48.4	✓ Active
2C:5A:0F:5D:B8:97	10.6.0.5	-	✓ Healthy	ifav22-leaf14	eth1/51.4	✓ Active
00:11:22:00:57:02	20.17.0.27 20:17::27	-	✓ Healthy	ifav22-leaf14	eth1/54.2.3907	✓ Active
F8:7A:41:47:C5:87	192.200.43.2	-	✓ Healthy	ifav22-spine9	eth1/53/4	✓ Active
2C:5A:0F:5D:B8:97	10.6.0.13	-	✓ Healthy	ifav22-leaf13	eth1/52.4	✓ Active

4. Locate the IP address for the endpoint and click on the IP address to launch the IP details page for the endpoint. For example, you can click on the endpoint **20.18.0.28**.

5. Click on the **Traffic analytics** tab.

IP details for IP 20.18.0.28

Refresh View in topology

Overview Endpoint history Anomalies **Traffic analytics** Trends and statistics Flow collections

Traffic analytics score reached Healthy ✔ Healthy

This score change generated 0 anomalies over the last 2 hours [View anomalies](#)

Services hosted on this endpoint

Filter by attributes

Service port	Conversation attributes ⓘ	Location	VLAN	Traffic analytics score ▾	Category	Protocol	TCP SYN Count (Intra-fabric)	TCP RST Count (Intra-fabric) ⚙
8018	External, Local	Local	-	✔ Healthy	Msite	TCP	6892	6892
8038	External, Local	Local	-	✔ Healthy	Udp_Msite	UDP	0	0

Connections to other services and IPs from this endpoint by Traffic analytics score ▾ over the last 2 hours

Filter by attributes

Start Flow collection ⬆

Connections to other services and IPs from this endpoint by Traffic analytics score ▾ over the last 2 hours

Filter by attributes

Start Flow collection ⬆

Endpoint	Conversation attributes ⓘ	Service port	Interface	Traffic analytics score ▾	Category	Protocol	Tenant	⚙
☐ 20.11.0.21	External	8031	eth1/7/1	✔ Healthy	Udp_Msite	UDP	-	te
☐ 20.11.0.21	External	8011	eth1/7/1	✔ Healthy	Msite	TCP	-	te
☐ 20.13.0.23	Local	8033	eth1/54/1	✔ Healthy	Udp_Msite	UDP	-	te
☐ 20.13.0.23	Local	8013	eth1/54/1	✔ Healthy	Msite	TCP	-	te



Service port numbers are displayed for TCP traffic; for ICMP traffic, no service port is displayed.

6. In **Service hosted on this endpoint** table, click on the service port to launch the **Service Endpoint Details** page for the endpoint.



The system does not display EPG, tenant, and VRF metadata for a remote external client in a different VRF instance. Choose the Nexus Dashboard site where the L3Out is defined to view this metadata.

Find traffic analytics conversations for an endpoint in a fabric group

Traffic analytics conversations and service endpoints can also be found from the fabric group page by navigating to the endpoint details, which filters all the traffic analytics records that we have for that specific endpoint.

1. Navigate to the **Manage > Fabrics > Fabric groups** page.
2. Click the *fabricgroup* name.
3. In the fabric group **Overview** page, navigate to **Connectivity > Endpoints**.

You can also navigate to the **Connectivity** area in your fabric group's **Overview** page and click on **Endpoints** to launch the **Endpoints** page.

← Back
tel7-nx-nx Refresh View in topology Actions ▾

Overview Inventory **Connectivity** Segmentation and security Configuration policies Anomalies Advisories History

Interfaces **Endpoints** Links Inter-fabric Virtual infrastructure Flows

🕒 Current ▾

Total 135 Anomaly level 1 134

Status == Active × Edit Clear all

MAC	IP addresses	Anomaly level	Connected to	VRF	Fabric	Tenant	
00:21:01:00:00:02		Warning	tel7-gx-leaf2	myvrf_50000	tel7-site-2	default	
00:11:01:00:00:2B		Healthy	tel7-fx-leaf3	myvrf_50000	tel7-site1-new	default	
00:11:01:00:00:13		Healthy	tel7-fx-leaf3	myvrf_50000	tel7-site1-new	default	
00:11:01:00:00:43		Healthy	tel7-fx-leaf3	myvrf_50000	tel7-site1-new	default	
00:11:01:00:00:5B		Healthy	tel7-fx-leaf3	myvrf_50000	tel7-site1-new	default	
00:11:01:00:00:3B		Healthy	tel7-fx-leaf3	myvrf_50000	tel7-site1-new	default	
00:11:01:00:00:33		Healthy	tel7-fx-leaf3	myvrf_50000	tel7-site1-new	default	
00:11:01:00:00:19		Healthy	tel7-fx-leaf3	myvrf_50000	tel7-site1-new	default	
00:11:01:00:00:0F		Healthy	tel7-fx-leaf3	myvrf_50000	tel7-site1-new	default	
00:11:01:00:00:4F		Healthy	tel7-fx-leaf3	myvrf_50000	tel7-site1-new	default	

135 items found Rows per page 10 ▾ < 1 2 3 4 5 ... 14 >

4. Locate the IP address for the endpoint and click on the IP address to launch the IP details page for the endpoint.
5. Click on the **Traffic analytics** tab.

Endpoint details for IP Refresh View in topology

[Overview](#)
[Endpoint history](#)
[Anomalies](#)
[Traffic analytics](#)
[Flows](#)

Traffic analytics score reached Warning Warning
 This score change generated 2 anomalies over the last 2 hours [View anomalies](#)

Services hosted on this endpoint
 Filter by attributes

Service port	Conversation attributes ⓘ	Location	VLAN	Security group ⓘ	Traffic analytics score ▲	Category ▼	Protocol	TCP SYN Count (Intra-fabric)	TCP RST Count (Intra-fabric)	Conversation Count (Intra-fabric)	ⓘ
5000	External	Local	467	13139	Warning	Analytics	TCP	0	0	0	

Connections to other services and IPs from this endpoint
 over the last 2 hours
 Filter by attributes

Endpoint	Conversation attributes ⓘ	Location	Security group ⓘ	Service port	Fabric name	Switch	Interface	Traffic analytics score ▲	Category ▼	ⓘ
 42.42.42.12	External	External	-	-	tel7-site1-new	tel7-fx-leaf3	eth1/2	Major	-	



Service port numbers are displayed for TCP traffic; for ICMP traffic, no service port is displayed.

- In **Service hosted on this endpoint** table, click on the service port to launch the **Service Endpoint Details** page for the endpoint.

Use the flow troubleshoot workflow

The flow troubleshoot workflow enables you to collect all the flow records between two endpoints. Nexus Dashboard allows you to specify the duration for flow collection and then collect records between specific endpoints for the specified duration. As a result you can view the path visualization, 5-tuple flow information, and any issues seen on individual flows.



The flow troubleshoot workflow feature is not supported in **Traffic analytics compatibility** mode.

Follow these steps to use the flow troubleshoot workflow.

- Navigate to **Analyze > Analyze Hub > Traffic analytics**.
- Choose a fabric from the drop-down list.
- Choose a time range from the drop-down list. By default the current time (last 2 hours) is chosen.
- Click on the **Services** tab and scroll down to the endpoints table.

5. Click on a port number listed in the **Service port** column in the endpoints table to view the **Service Endpoints Details** page.

1. On the **Service Endpoint Details** page, scroll down to the endpoints table.
2. Choose the radio button next to the Endpoint IP address and then click **Start Flow collection**. The **Duration for Flow collection** dialog box appears.

Service Endpoint Details for 202.1.1.18:90 Category: Udp_Drop

202.1.1.18	90	-	Jan 22 2026, 04:54:34.978 PM	-	-
VRF tenant1	Protocol UDP	Security group -	Switches -leaf4	Interfaces eth1/8/3	Fabric -

Top clients by Traffic analytics score

Filter by attributes

Warning

Healthy

Jan 22, 2026, 3:00 PM Jan 22, 2026, 3:27 PM Jan 22, 2026, 3:55 PM Jan 22, 2026, 4:22 PM Jan 22, 2026, 4:50 PM

Endpoint	Conversation attributes	Location	Fabric	Traffic analytics score	Switch	Interfac
☒ 23.0.0.55	Transit	External	-	Major	ifav22-leaf15	eth1/6/3

3. From the drop-down list, choose the duration for collecting flow records. You can choose a time range of 3,5,7,10, 15, or 20 minutes.
4. Click **Start** to begin the flow collection.

Service Endpoint Details for 202:1:1::18:90 Category: Udp_Drop

Jan 22 2026 03:24:26 PM - Jan 22 2026 05:24:26 PM

Overview Trends and statistics Anomalies **Flow collections**

Filter by attributes

Actions

Source	Source VRF	Destination	Destination port	Destination VRF	Protocol	Start time	End time	Collection status
☐ 23::55	tenant1	202:1:1::18	90	tenant1	UDP	Jan 22 2026, 5:25:22 PM	-	Scheduling

5. After the **Collection status** displays **Completed**, click **View records** to view the flow record details for that specific service endpoint.

Flow records between 23.0.0.55 and 202.1.1.18

Job details

Start time

Jan 22 2026
05:21:31.063 PM

End time

Jan 22 2026
05:31:47.746 PM

Collection status

Completed

Source address

23.0.0.55

Source VRF

tenant1

Destination address

202.1.1.18

Destination VRF

tenant1

Destination port

90

Protocol

UDP

Filter by attributes

Anomaly level	Record time	Switches	Protocol	Latency (µs)	Packet drop indicator	Source		
						Address	TCP/UDP port	Address
Major	Jan 22 2026 05:25:02.491 PM	ifav22-leaf4	UDP	0 µs	60	202.1.1.18	90	23.0.0.55
Major	Jan 22 2026 05:26:02.703 PM	ifav22-leaf4	UDP	0 µs	60	202.1.1.18	90	23.0.0.55
Major	Jan 22 2026 05:28:02.199 PM	ifav22-leaf4	UDP	0 µs	21	202.1.1.18	90	23.0.0.55
Major	Jan 22 2026 05:23:03.906 PM	ifav22-leaf4	UDP	0 µs	59	202.1.1.18	90	23.0.0.55
Major	Jan 22 2026 05:24:02.279 PM	ifav22-leaf4	UDP	0 µs	60	202.1.1.18	90	23.0.0.55
Major	Jan 22 2026 05:27:02.915 PM	ifav22-leaf4	UDP	0 µs	59	202.1.1.18	90	23.0.0.55
Healthy	Jan 22 2026 05:24:00.806 PM	ifav22-leaf15 View all (3 total)	UDP	3.174 µs	0	23.0.0.55	124	202.1.1.18
Healthy	Jan 22 2026 05:28:00.093 PM	ifav22-leaf15 View all (3 total)	UDP	3.174 µs	0	23.0.0.55	124	202.1.1.18
Healthy	Jan 22 2026 05:22:59.754 PM	ifav22-leaf15 View all (3 total)	UDP	3.174 µs	0	23.0.0.55	124	202.1.1.18
Healthy	Jan 22 2026 05:24:59.858 PM	ifav22-leaf15 View all (3 total)	UDP	3.174 µs	0	23.0.0.55	124	202.1.1.18

6. To view the flow collection for a fabric, navigate to **Manage > Fabrics**, choose a fabric, and click **Connectivity > Flow Collections**.
7. To perform flow collection for non-TCP flows follow these substeps.
 - a. Navigate to **Analyze > Analyze Hub > Traffic analytics**.
 - b. Choose a fabric from the drop-down list.
 - c. Choose a time range from the drop-down list. By default the current time (last 2 hours) is chosen.
 - d. Click on the **Endpoints** tab and scroll down to the endpoints table.
 - e. Choose the radio button next to the Endpoint IP address and then click **Start Flow collection**. The **Duration for Flow collection** dialog box appears.
 - f. From the drop-down list, choose the duration for collecting flow records. You can choose a time range of 3,5,7,10, 15, or 20 minutes.
 - g. Click **Start** to begin the flow collection.

- h. After the **Collection status** displays **Completed**, click **View records** to view the flow record details for that specific service endpoint.

Exploring the energy usage for your network using Energy Management

The Nexus Dashboard energy management report helps you monitor, predict, and improve your network's energy usage, its related carbon emissions, and its total energy cost. The energy management report enables you to get insights on energy utilization, CO2 emissions, and energy cost for all your fabrics 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 fabrics for every single day in the chosen month. This data is then combined with the Cisco Energy Manager to provide greater insight into what that usage means in terms of energy cost, estimated emissions, and estimated switch power consumption. For more information about the Cisco Energy Manager, see [Cisco Energy Manager](#).

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 fabrics' 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 or Panduit power distribution units (PDUs) are using.
- Estimated Emissions gives you insight into the impact your fabrics 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 the energy management report: one for switches only, and one for PDUs.

- Switch Data: Uses only the electricity data reported by individual switches added to a fabric.
- 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 the energy management report, you can:

- Better anticipate increases in your fabrics' energy bills so that your budgets more accurately reflect real-world usage.
- Better follow the hourly energy usage of an individual fabric. By spreading out usage to avoid peak hours surcharges, you may be able to lower your electricity bill over time.
- See the direct energy management 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 toward meeting ESG goals.



The retention time for the energy management report in Nexus Dashboard is 12

months.

Cisco Energy Manager

The Cisco Energy Manager is a service developed by Cisco that collects data from various data providers and consolidates the GHG emissions and the source of the energy from the data. The Cisco Energy Manager is hosted in a Cisco Intersight cloud.

View the energy management report for switches

1. Navigate to **Analyze > Analysis Hub > Energy management**.
2. Choose an online fabric or multiple online fabrics from the drop-down menu.
3. Choose a time range from the drop-down menu.
4. Use the **Display data from** toggle to display data from switches.
5. Click **Generate report**.

The energy management report displays At A Glance, Cost, Energy, and Emissions information for a particular fabric in the chosen month.

6. Examine the **At A Glance** area to see a summary of the estimated cost, estimated switch power consumption, and estimated emissions in the chosen month. Click the **Learn More** icon for more information.

Energy Management Actions

Showing latest data as of Jun 22nd 2026, 7:59 AM.
Emissions are estimated based on fabric locations and our [methodology](#). Before using Energy Management, you should review our [disclaimers](#).

All fabrics (7) This month display data from Switches PDU's


Using average emissions and cost values
To calculate cost estimates, we're using values based on the average cost of grid energy for each region. If you'd like, you can customize your average cost in [Fabric energy settings](#) for a more accurate estimate.

Energy management data might not yet be available.
Energy management data of fabrics that were onboarded within the last three hours may not yet be available.


June at a glance
Emissions are estimates based on fabric locations and utilities' self-reported energy sources, plus Cisco energy manager. You can learn more about our methodology [here](#)

Monthly summary		
Estimated cost	Estimated Switch power consumption	Estimated emissions
\$318.25	3182.53 kWh	395.59 kgCO2e

7. Examine the **Cost** area to see the estimated daily cost in the chosen month and share of daily cost per fabric.




- (Optional) From the **Actions** menu, choose **Fabric Energy Settings** to customize your average cost for the current month for a more accurate estimate. To calculate cost estimates, Nexus Dashboard uses values based on the average cost of grid energy for each region.
- Examine the **Energy** area to see the energy usage in the chosen month in kWh.



- Examine the **Emissions** area to see the total emissions or efficiency index per fabric, 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 chosen month.

For total emissions or efficiency index per fabric, use the toggle to view the information in graphical format or tabular format.



- Examine the **Top 5 devices** area to see the top 5 devices for the highest estimated cost, most energy consumed, and highest estimated greenhouse gas (GHG) emissions.

Top 5 devices
Showing devices by highest est. Cost, energy consumption, and est. GHG emissions for the selected time period


[View all devices](#)

Estimated cost	Estimated energy	Estimated emissions
ifav22-leaf4 \$31.93 USD estimated	ifav22-leaf4 319.39 kWh energy consumed	ifav22-spine7 49.59 kgco ₂ e
ifav22-leaf3 \$24.94 USD estimated	ifav22-leaf3 249.45 kWh energy consumed	ifav22-leaf11 43.11 kgco ₂ e
ifav22-leaf8 \$18.85 USD estimated	ifav22-leaf8 188.54 kWh energy consumed	ifav22-leaf12 41.33 kgco ₂ e
ifav22-leaf9 \$18.66 USD estimated	ifav22-leaf9 186.62 kWh energy consumed	ifav22-leaf14 37.61 kgco ₂ e
ifav22-spine1 \$18.17 USD estimated	ifav22-spine1 181.78 kWh energy consumed	ifav22-ipn 29.93 kgco ₂ e

Click **View all devices** to see the data for all devices, not just the top 5.

- Choose a fabric from the fabric drop-down 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 chosen month. The minimum period before you can generate the next report is 3 hours.




View the energy management report for PDUs


1. Navigate to **Analyze > Analysis Hub > Energy management**.
2. Choose an online fabric or multiple online fabrics from the drop-down menu.
3. Choose a time range from the drop-down menu.
4. Use the **Display data from** toggle to display data from PDUs.
5. Click **Prepare Report**.

The energy management report displays At A Glance, Cost, Energy, and Emissions information for a particular fabric in the chosen month.

6. Examine the **At A Glance** area to see a summary of the estimated cost, estimated switch power consumption, and estimated emissions in the chosen month. Click the **Learn More** icon for more information.



7. Examine the **Cost** area to see the estimated daily cost in the chosen month and share of daily cost per fabric.




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

9. Examine the **Energy** area to see the energy usage in the chosen month in kWh.



10. Examine the **Emissions** area to see the total emissions per fabric, 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 chosen month.

For total emissions per fabric, use the toggle to view the information in graphical format or tabular format.



11. Examine the **Top 5 Devices** area to see the top 5 devices for the highest estimated cost, most energy consumed, and highest estimated greenhouse gas (GHG) emissions.

Click **View all devices** to see the data for all devices, not just the top 5.

12. Choose a fabric from the fabric drop-down 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 chosen month. The minimum period before you can generate the next report is 3 hour.

Hourly Energy Mix

Amount of total energy used during each three-hour reporting period by source, over all of the days for December

- Solar

Nuclear

Biomass

Oil

Hydro

Storage

Unknown
- Wind


Coal

Hydro

Battery

Storage

Gas



Comparing the configurations between two time periods for your fabric using Delta analysis

Nexus Dashboard performs analysis of fabrics 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 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. Delta analysis includes microsegmentation endpoint groups (uSeg EPGs) in the EPG resource view on the **Health Delta** page. You can compare the health and anomaly status of uSeg EPGs between two snapshots.

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

Policy Delta for ACI

Policy Delta for ACI fabrics analyzes the differences in the policy between the two snapshots and provides a co-related view of what has changed in the ACI Fabric.

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

Policy Delta for standalone NX-OS

Policy Delta for standalone NX-OS fabrics analyzes the changed nodes or switches across two snapshots and obtains a co-related view of what has changed in the NX-OS switches.

See [View policy delta analysis](#) for more details.

Guidelines and limitations for delta analysis

- For NX-OS fabrics, 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 fabric 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.

- Delta Analysis might not associate all active anomalies with their affected resources. As a result, an affected resource could appear healthy in the **Health Delta** view even when active anomalies exist. Additionally, when selecting an unhealthy resource, the **All Anomalies** table may not display all active anomalies related to that resource.
- For ACI fabrics, the **APIC Configuration Export Policy** must be of the same format (XML/JSON) for both the snapshots.
- For ACI fabrics, the policy delta will not be performed if there are any APIC configuration export policy collection errors.
- For ACI fabrics, the toggle for 'Include Acknowledged Anomalies' which allows you to filter out the acknowledged anomalies from the results displayed does not show anomalies that are manually acknowledged.
- For NX-OS fabrics, when you choose a switch in the **Changed Nodes** area, in **Policy Delta**, the difference in the configuration between the two snapshots is displayed.
- For NX-OS fabrics, for Policy Delta, **Audit Logs** is not currently supported.

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**. This is not a requirement for NX-OS fabrics.

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 **Fabric** to choose the fabric.
3. Click **Choose Earlier Snapshot** and choose the first snapshot for the delta analysis. Click **Apply**.
4. Click **Choose Later Snapshot** and choose the second snapshot for the delta analysis. Click **Apply**.



The two snapshots chosen for the delta analysis must belong to the same fabric.

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, you can **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 another delta analysis.

1. (Optional) From the Status column, choose 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 page of the rule creation as a banner.

View delta analysis

The delta analysis 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. Click any delta analysis to view more details.

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
- Fabric
- 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:

The toggle for 'Include Acknowledged Anomalies' allows you to filter out the acknowledged anomalies from the results displayed if enabled. If it is disabled, manually acknowledged anomalies are included in the Anomaly Count.

- **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.

Delta analysis now performs an object delta rather than a count delta. So along with the count, you can now view how many anomalies were cleared, how many are unchanged and how many are new anomalies.

The anomaly count also displays the difference for the different types of anomalies. It is displayed for **Critical**, **Major** and **Warning**.

- **Delta by Resources:** Displays the count of resources by type that are new, lost or unchanged. You can also specifically view the resources whose count has changed by clicking the **View Changed Only** toggle. The filter bar allows you to filter the data by resource. The gear icon allows you to customize the columns as per your view. The table also shows the count delta and the health delta. Count delta includes both healthy and unhealthy resources. Healthy resources will not have any anomalies associated with them if filtered. Health delta shows only unhealthy resources and will return anomalies if filtered by anomalies.

If you click any of the counts for the Count and Health Delta resources, you can view the list of resources along with the node information. For SVI, the node column provides the VNI association.

- **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.

For ACI fabrics, the Anomalies can be listed for the following kinds of snapshots:

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

For NX-OS fabrics, the Anomalies can be listed for the following types:

- New
- Unchanged
- Cleared
- From Earlier Snapshot
- From Later Snapshot

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.

For ACI fabrics, you can filter the results based on the following attributes:

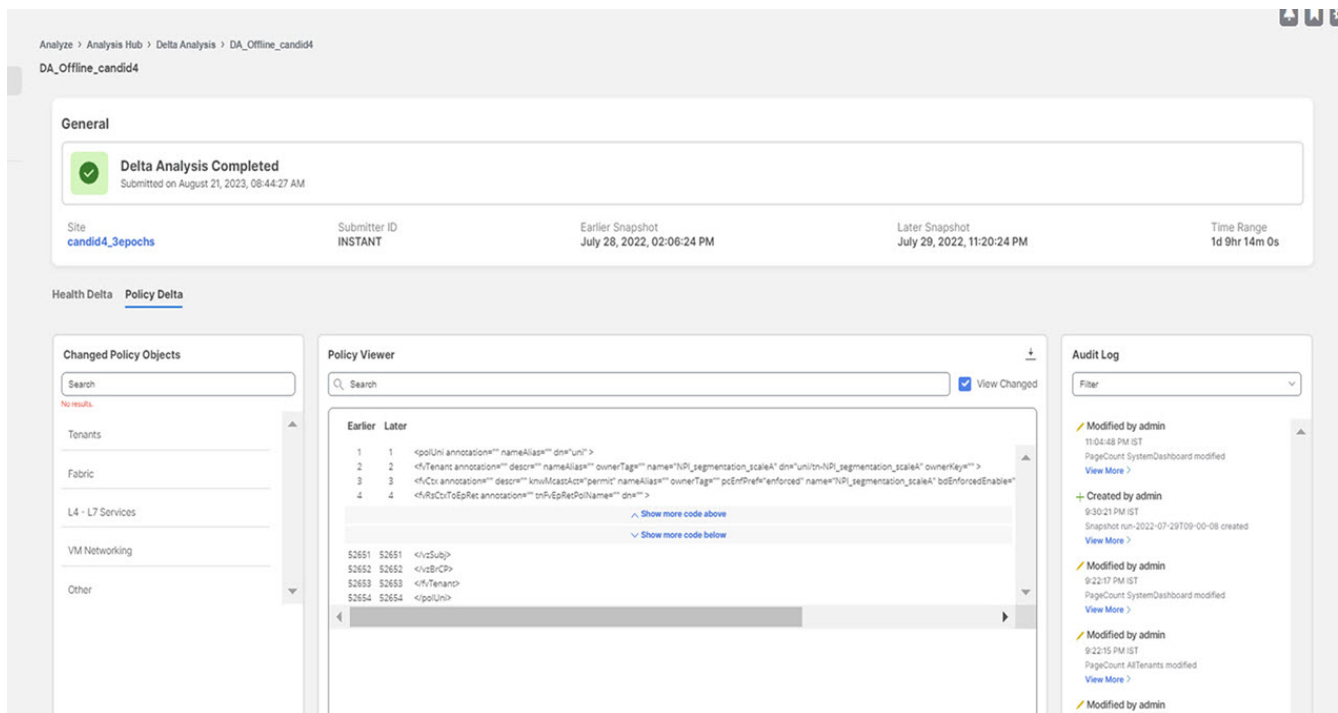
- Anomaly Level
- App Profile DN
- BD DN
- Title
- Contract DN
- EPGs
- External Routes
- Interfaces
- Internal Subnets
- L3Out DN
- Leaf DN
- Tenant DN
- Endpoints
- VRF DN

For NX-OS fabrics, you can filter the results based on the following attributes:

- Border Gateways (Leaf)
- Border Leafs (Leaf)
- Interfaces
- L2VNIs
- L3VNIs
- Leafs
- Spines
- SVIs
- VLANs
- VNIs
- VPCs
- VRFs

Choose an anomaly to view the anomaly details.

View policy delta analysis



Click **Policy Delta** to view the policy changes across the two snapshots. Policy Delta includes Changed Policy Object and Policy Viewer. ACI fabrics also include Audit Log.

1. **Changed Policy Object** 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. Choose 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. Nexus Dashboard 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 choose 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.

Viewing the impact of configuration changes using pre-change analysis (ACI fabrics)

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

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

Analyze > Analysis Hub > Pre-Change Analysis

Pre-Change Analysis

Create New Pre-Change Analysis

Filter by attributes

☐ Pre-Change Analysis Name	Assurance Entity Name	Base Epoch	Analysis Status	Analysis Submission Time	Submitter ID	
☐ test2	ACI-Paris	February 01, 2024, 11:24:35 PM	✓ Completed	February 01, 2024, 11:48:29 PM	Local: admin	...
☐ test	ACI-Paris	February 01, 2024, 11:24:35 PM	✓ Completed	February 01, 2024, 11:43:47 PM	Local: admin	...

2 items found

Rows per page 10 < 1 >

© 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 chosen 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

Analyze > Analysis Hub > Pre-Change Analysis > test2

test2

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

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:

- Fabric 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 [Comparing the configurations between two time periods for your fabric using Delta analysis](#).

To understand the data displayed under Compliance Analysis, view [Defining communication and configuration rules using Compliance \(ACI fabrics\)](#).

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 chosen 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, consumed 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.

Guidelines and limitations for pre-change analysis

When using pre-change analysis follow these guidelines and limitations:

- Pre-change analysis can be conducted for fabrics and uploaded config 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 chosen.
- 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. Nexus Dashboard 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.
- You can upload a JSON or an XML Cisco APIC configuration file to run a pre-change analysis job.
 - The maximum file size is 10 MB for vND and 50 MB for pND.
 - When Nexus Dashboard extracts a tenant configuration from Cisco APIC during a Pre-Change Analysis job, the extracted tenant configuration must not exceed 10 MB for vND or 50 MB for pND. If the tenant configuration exceeds this size, the pre-change analysis job fails.
 - 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 fabric, you can only run a pre-change analysis job one at a time.
- If you modify an object that belongs to a tenant, the pre-change analysis file size for that tenant cannot be more than 10 MB.

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 files 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 Nexus Dashboard, 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, Nexus Dashboard 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 Nexus Dashboard.
- Pre-change analysis is not supported in Cisco ACI configurations that contain service chains.
- Nexus Dashboard 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 the 3.2(x) release, 3.2(9h) and earlier
 - For the 4.0(x) release, 4.0(1h) and earlier
 - For the 4.1(x) release, 4.1(2x) and earlier
 - For the 4.2(x) release, 4.2(7s) and earlier
 - For the 5.0(x) release, 5.0(2e) and earlier
 - For the 5.1(x) release, 5.1(4c) and earlier
 - For the 5.2(x) release, 5.2(4d) and earlier
 - For the 5.3(x) release, 5.3(1b) and earlier
 - For the 6.0(x) release, 6.0(4c) and earlier
 - For the 6.1(x) release, 6.1(4h) and later
 - For the 6.2(x) release, 6.2(1g) and later

Create pre-change analysis job

1. Navigate to **Analyze > Analysis Hub > Pre-change**.
2. In **Pre-Change Analysis**, click **Create new pre-change analysis**. In **Create new 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 **Fabric** field, choose the appropriate fabric.
- d. In the **Snapshot** field, specify the appropriate snapshot.

Change simulation

- a. Under **Change simulation**, 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, choose 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** 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.

Collecting and analyzing your device logs using Log Collector

The Log Collector feature enables you to collect support logs for local usage and optionally 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 fabric 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 fabric 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 uses the device connectivity issue notifier on 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 notifies the following:

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

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

Guidelines and Limitations

- If you see the error message **Unable to Add Fabric** when you try to create a new log collector job, this can sometimes occur if the local time is not synchronized properly on the PC where you are accessing the Nexus Dashboard GUI. Synchronize the local time on your PC and attempt to create the log collector job again to resolve the issue.

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 fabric and displays the latest log collections.

The filter bar allows you to filter the logs by Status, Name, Node, start time, and end time. For standalone NX-OS fabrics, you can also filter by Type.

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.

The page also displays the log collection jobs in a tabular format. The jobs are sorted by status. Choose 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. Choose the log collection job in the table to display additional details. The **Log Collection** status displays information such as status, general information, and node details.



You can save TAC assist job details as a PDF with the browser print option (Only supported on Chrome and Firefox).

Upload logs to Cisco Intersight Cloud

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

Choose **Analyze > Analysis Hub > Log Collector > Create Log Collector Job**.

1. In the Log details section, enter a name in the **Name** field.
2. Click the drop-down list in the **Fabric name** field to choose a fabric.
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. In the **Select Nodes** section, use the **Switches** and **Controllers** tabs to choose the required nodes.
6. Click **Next**. The Summary section displays the entire job summary.
7. Click **Submit** to initiate the log collection process. When the job is completed, Nexus Dashboard displays a message to confirm the job was successfully completed.
8. Click **View log collector jobs** to view log collection information in the Log Collector table.
9. Click on a job name listed in the Name column to view the details.
10. Choose the node and click  icon.
11. Click **Upload File to TAC Assist** to upload a single file for the chosen node manually.
12. Click **Upload All** to upload all the log files generated for the job.

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 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.

Collecting information on bugs that might affect your network using Bug scan

Nexus Dashboard collects technical support information from all the devices and runs them against known set of signatures, and flags the corresponding defects and PSIRTs. Nexus Dashboard also generates advisories for PSIRTs and anomalies for defects. See [Identifying Advisories in Your Nexus Dashboard](#) and [Detecting Anomalies in Your Nexus Dashboard](#) to learn more about metadata support. Nexus Dashboard bug signatures are updated through metadata information, see the "Metadata" section in [Working with System Settings](#) information. The metadata file is updated automatically if Nexus Dashboard is connected to [Cisco Intersight](#) or can be manually updated. For more information, see the [Enhancements to Bug scan](#) section.

The Bug scan feature collects technical support logs from devices in a fabric and scans them for bugs that could have been hit. If the CPU and memory usage of the devices is below the set threshold of 65% then the tech support logs are collected and the Bug scan is carried out for the devices. If the CPU and memory usage is above the set threshold, the devices are excluded from the Bug scan and eventually will be reconsidered for the next default Bug scan or when you run an on-demand Bug scan for that device.



CPU memory metrics for the devices must be available for Bug scan jobs to be scheduled.

The switch connectivity must be healthy on the device for the Bug scan process to successfully collect logs from the device.

You can also run an on-demand Bug scan for a fabric. For more information, see the On-Demand Analysis section in [Understanding your Nexus Dashboard Fabrics](#).

Enhancements to Bug scan

Beginning with Nexus Dashboard 4.2.1, the Bug scan feature is enhanced to classify bugs as Active (in addition to Known) based on device software version, running configurations, and log/service analysis. Bug scan behavior is changed in Nexus Dashboard 4.2.1, particularly in its interaction with Cisco Intersight and metadata management. When your Nexus Dashboard is initially deployed and not connected to Cisco Intersight, the Bug scan feature will not work. As a result, detection of Active bugs, PSIRTs, and Known bugs is unavailable. The pre-packaged subset of metadata includes Advisories, Release information, and Best practices. The pre-packaged metadata does not contain comprehensive bug data, signatures, Field Notices, Product Security Incident Response Team (PSIRT) information, or Known bug details. Once metadata is downloaded, Bug scan functionality becomes available, and you can obtain detailed bug data, signatures, PSIRTs, and Known bug details. For more information, see [Enable Bug scan](#).



Nexus Dashboard UI will display a bell icon notification to alert the user to update metadata. The alert notification states, "The Bug scan feature will not be available until Nexus Dashboard metadata has been updated. You can update the metadata. This can be done by configuring the Intersight Device Connector, or by manually uploading a metadata file under **System Settings > Metadata**. For more information click below". You can click on the link to follow the update instructions. Similarly, the **Summary** area on the **Analyze > Analysis Hub > Bug scan** page will display a

notification alerting the user to update the metadata. You can configure the **Intersight Device Connector** under **Admin > Intersight**, or by manually uploading a metadata file under **Admin > System Settings**. For more information on generating metadata files for non-intersight connected Nexus Dashboard Clusters, please refer [here](#).



If metadata is not downloaded, the **Known bugs** page displays "No known bugs, No active bugs." You can only access a limited set of **Advanced Advisories** that rely on running switch commands. Features that require an internet connection, such as serial number validation, will not be available.

Bug scan schedule

Bug scan runs for all fabrics onboarded to Nexus Dashboard using a timer-based mechanism. Once a fabric is onboarded, all switches in that fabric are included for Bug scan processing. This behavior is system-defined and not user-configurable.



Bug scan is enabled by default, you cannot disable the Bug scan feature.

Bug scan enforces a minimum 7-day gap per device between scans, regardless of whether a previous scan completed successfully or failed to execute (for example, due to high CPU or memory utilization on the switch). If a scan does not execute or fails, it will be attempted again when the timer triggers, subject to the same 7-day constraint. In addition to timer-based execution, on-demand Bug scans or device reloads may trigger a scan earlier than the 7-day interval. When this occurs, the 7-day window is reset and starts from the time of that execution. It is expected that a Bug scan will be attempted on a device approximately every 7 to 14 days, depending on the number of nodes and fabrics in the Nexus Dashboard cluster.

On-demand Bug scans are a user-triggered collection of technical support information from selected switches of a fabric. Those on-demand scans are prioritized over auto-scheduled runs and do not consider the CPU and memory metrics. On demand bug scans can be executed for up to 10 switches concurrently. If auto-scheduled Bug scan is in progress and on-demand Bug scan is initiated on other switches, based on the available resources in the Nexus Dashboard nodes the on-demand Bug scan will start while the current Bug scan is in progress or after the current Bug scan is completed.

Only one on-demand bug scan job can be triggered in {PlatformFullName}.



A Bug scan is triggered automatically in the following scenarios:

- Switch or APIC controller upgrade or downgrade
- Switch or APIC controller reloads

Guidelines and limitations

- You must either reconnect to Intersight to download metadata or apply air-gapped metadata in these scenarios:
 - When a Nexus Dashboard node restarts or reloads, you might see a banner indicating that

newer metadata is available. In this case, you must either reconnect to Intersight or apply air-gapped metadata.

- After performing a backup and restore operation, you might see a banner indicating that newer metadata is available. You must then either reconnect to Intersight or apply air-gapped metadata.
- If one of the metadata update pods restarts, you may see a banner indicating that newer metadata is available. The user must either reconnect to Intersight or apply air-gapped metadata.

Enable Bug scan

Before you begin

- Ensure you have an active Cisco Intersight account.

Once the Bug scan (full functionality) is enabled, you can download the complete metadata package which includes:

- Comprehensive bug data
- Signatures
- Field Notices
- PSIRT information
- Known bug details

Nexus Dashboard offers two options to enable Bug scan, you can choose one of these options to enable Bug scan:

- [Enable Bug scan using Cisco Intersight account](#)
- [Upload metadata using air-gapped mechanism](#)

Enable Bug scan using Cisco Intersight account

Follow these steps to enable Bug Scan using Cisco Intersight account.

1. Navigate to the **Admin > Intersight** page. If you are already connected to Intersight, the **Intersight Device Connector** page will display your device ID and account details under the **Device ID** and **Claimed to Account** fields.
2. Log into your Cisco Intersight account and navigate to **System > Targets** and click on **Claim a New Target**.
3. On the **Select Target Type** page, click on **Cisco Nexus Dashboard** under the **Network** category.
4. Enter the values for **Device ID** and **Claim Code** on the **Claim a new target** page.
5. Click on **Claim** and wait for a few seconds. Once claimed, your Nexus Dashboard **Intersight Device Connector** page displays a **Claimed** status. Your device Id and the Intersight account details will also display under the **Device ID** and **Claimed to Account** fields.

After connecting to Intersight, Nexus Dashboard will automatically pull for updates every 4 hours. Once the full metadata is successfully downloaded, all Bug Scan features are enabled. Bug Scan jobs will begin to be scheduled automatically, and the notification regarding missing metadata will also


disappear.

Upload metadata using air-gapped mechanism

You can use the air-gapped mechanism to manually upload your metadata file when Cisco Intersight is not available.

Follow these steps to upload metadata using the air-gapped mechanism.

1. Navigate to **Admin > System Settings > Metadata**.
2. Click **Edit** in the **Metadata** card to view the **Generate key**.
3. Click **Generate key** to generate an encryption key that you can use to download the encrypted metadata file.
4. If you have not yet created an Intersight PVAPP account, you first need to create one at <https://intersight.com/pvapp>. You will be prompted to set up an account name. You can use 'NexusDashboard-Metadata' or a similar name to easily identify it from other Intersight accounts you may already have.
5. Login to Cisco Intersight account using the previously created 'NexusDashboard-Metadata' or an appropriate account.
6. Click the **Nexus Dashboard Metadata** tab on the **Software Download** page.



- o Click the download icon next to the latest metadata file available for download.
7. Copy the key you generated in Step 3 and paste the key in the **Download Nexus Patch Bundle** box and click **Download**.
 8. Wait for the download to complete.
 9. In Nexus Dashboard, navigate to **Admin > System Settings > Metadata** page.
 10. Upload the downloaded metadata file to the **Update metadata version** field.

View known bugs

The Bug Scan feature collects technical support logs from devices in a fabric and scans them for


bugs. You can view the known bugs affecting your network after the Bug Scan is completed.

Known Bugs - Bugs present in the software version that may potentially impact your network. These bugs do not require a bug signature but are based only on the software version match and not on switch configurations or tech-support.



Only Severity 1 to Severity 3 bugs attached to the Release Notes of a certain ACI or NX-OS release are considered known bugs.

1. Navigate to **Analyze > Analysis Hub > Bug Scan**.
2. Choose an online fabric or multiple online fabrics from the drop-down list.
3. Choose the software version from the drop-down list. The known bugs for the chosen fabrics and software versions are displayed.



4. The Summary area displays the overall bugs by severity. You can also view the Bugs per fabric or software version using the drop-down list.
5. In the Bugs area, use the filter bar to filter the bugs by bug ID, description, severity level, type, and affected nodes.
6. View the **Severity Level** donut chart to see the total number of bugs of Critical, Major, and Warning severity.
7. View the bugs table to see the filtered bugs.
 - a. Click the column heading to sort the bugs in the table.
 - b. Click the gear icon to configure the columns in the table.
 - c. Click **Bug ID** to view bug details.
8. Click **Run Bug Scan** to run an on-demand Bug Scan. Choose a fabric and click **Run Now**. For more information, see the On-demand analysis section in [Understanding your Nexus Dashboard](#)

View known bugs for an individual fabric

In Nexus Dashboard, you can also view the bugs for an individual fabric in the following ways:

1. Navigate to **Manage > Fabrics**
2. Choose **Online Fabrics** from the drop-down list.

Name	Anomaly Level	Advisory Level	Type	Connectivity to Nexus Dashboard Insights	Software Version ⓘ	Creation Time on Nexus Dashboard	
DC-ute11	Major	Warning	ACI	OK	6.0(5h)	May 13, 2024, 11:15:30 AM	...

3. In the Software Version column, click on a version to open the drawer that lists the software versions, number of nodes on each version, and total bugs.
4. Click **View all** in the drawer to view the known bugs for that fabric and version.

OR

1. Navigate to **Manage > Fabrics**.
2. Choose a fabric.
3. In the General area, click on **Show more** under Fabric bugs to open the software version drawer.
4. Click **View all** in the drawer to view the known bugs for that fabric and version.
 - For ACI fabrics, you can also click the APIC software version to directly view the known bugs for that version.

General
Showing most recently available data

Type
ACI

Connectivity to Nexus Dashboard Insights
OK

Telemetry Collection Status
Not Available

Software Version
6.0(5h)

Creation Time on Nexus Dashboard
May 13, 2024, 11:15:30 AM

Nexus Dashboard Insights Collector Configuration
IPv4

Telemetry Streaming Network
In-Band

Inventory
Showing most recently available data

Controllers
1

Switches
6

[View Hardware Resources](#)

[View Capacity](#)

Connectivity

Endpoints
14

L3 Neighbors
26

OR

1. Navigate to **Manage > Inventory**
2. Choose **Online Fabrics** from the drop-down list.
3. In the **Controllers** tab click the software version in the Controller table to directly view the known

bugs for that fabric and version.

4. Click the **Switches** tab and click the software version in the Switches table to directly view the known bugs for that fabric and version.
5. Click **Run Bug Scan** from the **Actions** drop-down list in any of these views to run an on-demand Bug Scan. For more information, see the On-demand analysis section in [Understanding your Nexus Dashboard Fabrics](#).



The option to view bugs through Manage > Fabric Software Management > Firmware Summary is no longer valid.

Tracking endpoints within a data center using Endpoint Locator (NX-OS fabrics)

The Endpoint Locator (EPL) feature allows real-time tracking of endpoints within a data center. The tracking includes tracing the network life history of an endpoint and getting insights into the trends that are associated with endpoint additions, removals, moves, and so on. An endpoint is anything with at least one IP address (IPv4 and/or IPv6) and MAC address. EPL feature is also capable of displaying MAC-Only endpoints. By default, MAC-Only endpoints are not displayed. An endpoint can be a virtual machine (VM), container, bare-metal server, service appliance and so on.



- EPL is supported for VXLAN BGP EVPN fabric deployments only in the Nexus Dashboard LAN fabric installation mode. The VXLAN BGP EVPN fabric can be deployed as Easy fabric, Easy eBGP fabric, or an External fabric (managed or monitored mode). EPL is not supported for 3-tier access-aggregation-core based network deployments.
- EPL displays endpoints that have at least one IP address (IPv4 and/or IPv6). EPL is also capable of displaying MAC-Only endpoints. Select the **Process MAC-Only Advertisements** checkbox while configuring EPL to enable processing of EVPN Route-type 2 advertisements having a MAC address only. L2VNI:MAC is the unique endpoint identifier for all such endpoints. EPL can now track endpoints in Layer-2 only network deployments where the Layer-3 gateway is on a firewall, load-balancer, or other such nodes.

EPL relies on BGP updates to track endpoint information. Hence, typically the Nexus Dashboard must peer with the BGP Route-Reflector (RR) to get these updates. For this purpose, IP reachability from the Nexus Dashboard to the RR is required. This can be achieved over in-band network connection to the Nexus Dashboard Data Network interface. There is no option to configure static routes for pods on ND, so the selected RRs must be reachable through the default data network gateway.

Some key highlights of the Endpoint Locator are:

- Support for dual-homed and dual-stacked (IPv4 + IPv6) endpoints
- Support for up to two BGP Route Reflectors or Route Servers
- Support real-time and historical search for all endpoints across various search filters such as VRF, Network, Layer-2 VNI, Layer-3 VNI, Switch, IP, MAC, port, VLAN, and so on.
- Support for real-time and historical dashboards for insights such as endpoint lifetime, network, endpoint, VRF daily views, and operational heat map.
- Support for iBGP and eBGP based VXLAN EVPN fabrics. The fabrics may be created as Easy Fabrics or External Fabrics. EPL can be enabled with an option to automatically configure the spine or RRs with the appropriate BGP configuration.
- You can enable the EPL feature for upto 4 fabrics.
- EPL is supported on Multi-Site Domain (MSD).
- IPv6 underlay is not supported.
- Support for high availability
- Support for endpoint data that is stored for up to 60 days, amounting to a maximum of 2 GB storage space.

- Support for optional flush of the endpoint data to start afresh.
- Supported scale: Maximum of 50K unique endpoints per fabric. A maximum of 4 fabrics is supported. However, the maximum total number of endpoints across all fabrics should not exceed 100K.

If the total number of endpoints across all fabrics exceeds 100K, an alarm is generated and is listed under the **Alarms** icon at the top right of the page. This icon starts flashing whenever a new alarm is generated.

- Persistent or External IP addresses are required to enable EPL. For each VXLAN fabric, a specific container is spawned running a BGP instance to peer with the spines of the fabric. This container must have a persistent IP associated that is then configured as a iBGP neighbor on the spines. A different container is used for each fabric, so the number of fabrics that are managed by Nexus Dashboard where EPL is enabled decides how many persistent IP addresses must be distributed for EPL. Also, the EPL establishes iBGP sessions only over the Nexus Dashboard Data interface. The reachability from Nexus Dashboard to the switches must be over the in-band, also known as the front-panel ports.
- You can disable promiscuous mode on the port-groups that are associated with the Nexus Dashboard Management or Data vNICs. The Persistent IP addresses are given to the pods (for example, SNMP Trap/Syslog receiver, Endpoint Locator instance per Fabric, SAN Insights receiver, and so on). Every POD in Kubernetes can have multiple virtual interfaces. Specifically for IP stickiness an extra virtual interface is associated with the POD that is allocated an appropriate free IP in the external service IP pool. The vNIC of the POD that has the Persistent IP shares the same MAC address of Nexus Dashboard bond0 or bond1 interface. Therefore, the POD sources the packets using the same MAC address of Nexus Dashboard bond0 or bond1 interfaces that are known by the VMware ESXi system.

If you are using a Virtual Cisco Nexus Dashboard Cluster before you begin, ensure that the Persistent IP addresses, EPL feature, and required settings are enabled. See the *Nexus Dashboard Installation and Upgrade Guide* for more information. === EPL behavior with NX-API configuration

When the **NX-API** option is enabled on supported fabrics, such as the VXLAN, Enhanced Classic LAN, eBGP, and Campus fabrics, the EPL feature behavior is as follows:

- Dashboard counter for dual attached endpoints is populated.
- Endpoints discovered under a vPC pair will have the switch name set to the specific switch the endpoint is found under.
- Endpoints under ports such as the vPC Peer-Link or as a port name that contains sup, for example, sup eth1(R), will be filtered out from being discovered.

For more information on **NX-API** field related information, see the Advanced section in [Fabric Management](#).

When the **NX-API** option is disabled, the EPL feature behavior is as follows:

- Dashboard counter for Dual Attached Endpoints is always 0.
- VLAN and port for endpoints are left empty.
- The VRF field for endpoints is set to VRF ID instead of VRF name.
- Endpoints under a vPC pair will have the switch name set to the name of both switches in the vPC pair.

- Endpoints with ports such as vPC Peer-Link(R) or has a port name that contains sup, for example, supeth1(R), will be discovered without the port name and displayed as part of the dashboard data.

Backup and Restore

EPL only backups data for fabrics that EPL has been configured. If EPL is disabled for a fabric (even if EPL has previously been configured there), then you cannot backup the data for that fabric. Also, you can backup only historical data (data on the Endpoint Search page).

If a backup is initiated when EPL is enabled, then when restoring the backup, the same external data IPs that EPL was using must be available on ND. If those IPs are not available, then select the **Ignore External Service IP Configuration** option in the restore backup form. However, there are chances that the EPL pods will be brought up with different IPs, so any existing EPL policies become invalid. If EPL was previously configured with the **Configure My Fabric** option, you need to disable and enable EPL so that the old policy is cleaned up and an updated policy is deployed. If you did not use the **Configure My Fabric** option, then manually update their config with the new IPs.

EPL Connectivity Options

Sample topologies for the various EPL connectivity options are as given below.

Nexus Dashboard Cluster Mode: Physical Server to VM Mapping

Refer to the *Nexus Dashboard Verified Scalability Guide* for more information.

Configuring Endpoint Locator

The Nexus Dashboard OVA or the ISO installation comes with two interfaces:

- Management
- Data

The Management interface provides reachability to the devices via the mgmt0 interface either Layer-2 or Layer-3 adjacent. This allows Nexus Dashboard to manage and monitor these devices including POAP. EPL requires BGP peering between the Nexus Dashboard and the Route-Reflector. Since the BGP process on Nexus devices typically runs on the default VRF, in-band IP connectivity from the Nexus Dashboard to the fabric is required. The data network interface can be configured during Nexus Dashboard installation. You can't modify the configured in-band network configurations.



The setup of Data network interface on the Nexus Dashboard is a prerequisite of any application that requires the in-band connectivity to the devices within fabric. This includes EPL and Network Insights Resources (NIR).

As mentioned earlier, for EPL to function, IP reachability must be present over the ND data interface to the spine BGP loopback interface IPs. Typically the reachability to the remote ND cluster is established via the border switches in the default VRF.

For redundancy purposes, it is always advisable to have the server on which the Nexus Dashboard is installed to be dual-homed or dual-attached. With the OVA Nexus Dashboard deployment, the server

can be connected to the switches via a port-channel. This provides link-level redundancy. To also have node-level redundancy on the network side, the server may be attached to a vPC pair of Leaf switches. In this scenario, the switches must be configured such that the HSRP VIP serves as the default gateway of the Data Network interface on the Nexus Dashboard.

After you establish the in-band connectivity between the physical or virtual Nexus Dashboard and the fabric, you can establish BGP peering.

During the EPL configuration, the route reflectors (RRs) are configured to accept Nexus Dashboard as a BGP peer. During the same configuration, the Nexus Dashboard is also configured by adding routes to the BGP loopback IP on the spines/RRs via the Data Network Interface gateway.



Cisco Nexus Dashboard queries the BGP RR to glean information for establishment of the peering, such as ASN, RR, and IP.

Configure Endpoint Locator from Nexus Dashboard user interface

Follow these steps to configure Endpoint Locator from Nexus Dashboard user interface.

- In the fabric **Overview** page, choose **Actions > Configuration > Configure Endpoint Locator**.
- Similarly, you can configure EPL on the **Topology** page. Right-click on the required fabric, then click **More > Configure Endpoint Locator**. The **Endpoint Locator** page appears.
- You can enable EPL for one fabric at a time. Select the switches on the fabric hosting the RRs from the drop-down list. Nexus Dashboard will peer with the RRs.

By default, the **Configure My Fabric** option is selected. This option only configures EPL as a BGP neighbor of the switch and this option does not configure network reachability between EPL and the switch. This knob controls whether BGP configuration will be pushed to the selected spines/RRs as part of the enablement of the EPL feature. If the spine/RR needs to be configured manually with a custom policy for the EPL BGP neighborhood, then this option should be unchecked. For external fabrics that are only monitored and not configured by Nexus Dashboard, this option is greyed out as these fabrics are not configured by Nexus Dashboard.

Select the **Process MAC-Only Advertisements** option to enable processing of MAC-Only advertisements while configuring the EPL feature.



If EPL is enabled on a fabric with or without selecting the **Process Mac-Only Advertisements** checkbox and you want to toggle this selection later, then you have to first disable EPL and then click **Database Clean-up** to delete endpoint data before re-enabling EPL with the desired **Process Mac-Only Advertisements** setting.

Select **Yes** under **Collect Additional Information** to enable collection of additional information such as PORT, VLAN, VRF etc. while enabling the EPL feature. To gather additional information, NX-API must be supported and enabled on the switches, ToRs, and leafs. If the **No** option is selected, this information will not be collected and reported by EPL.



For all fabrics except external fabrics, NX-API is enabled by default. For external fabrics, you have to enable NX-API in the external fabric settings by selecting the **Enable NX-API** checkbox in the **Advanced** tab of the External_Fabric_11_1 fabric

template.

Click the **Config preview** button to view the configuration that will be pushed to the switches while enabling EPL. You can copy and paste this configuration on spines or border gateway devices to enable EPL on external monitored fabrics.

Once the appropriate selections are made and various inputs have been reviewed, click **Submit** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled.

The Nexus Dashboard Data Service IP is used as BGP EVPN neighbor.

When the Endpoint Locator feature is enabled, there are a number of steps that occur in the background. Nexus Dashboard contacts the selected RRs and determines the ASN. It also determines the interface IP that is bound to the BGP process. Also, appropriate BGP neighbor statements are added on the RRs or spines in case of eBGP underlay, to get them ready to accept the BGP connection that will be initiated from the Nexus Dashboard. The external Nexus Dashboard Data Service IP address that is assigned to the EPL pod will be added as the BGP neighbor. Once EPL is successfully enabled, the user is automatically redirected to the EPL dashboard that depicts operational and exploratory insights into the endpoints that are present in the fabric.


For more information about the EPL dashboard, see [Monitoring Endpoint Locator](#).

Configuring Endpoint Locator for Single VXLAN EVPN Site

Before you begin:

In the following figure, the Nexus Dashboard service application is attached to the VPC pair of Leaf switches as it provides the link and node-level redundancy. The BGP instance running on EPL container establishes iBGP peering with the fabric spines. The iBGP peering is between Spine loopback addresses (loopback0) and EPL container persistent IP addresses. The loopback0 address of Spines is reachable via VXLAN Underlay, therefore, EPL container IP must have IP reachability towards the spines. We can configure an SVI on Leaf switches that can provide IP connectivity. The SVI will be a non-VXLAN enabled VLAN and will only participate in the underlay.

VXLAN EVPN Fabric-1



Follow these steps to configure endpoint locator for single VXLAN EVPN site.

1. You must configure persistent IP addresses on Cisco Nexus Dashboard. On Nexus Dashboard, choose **Admin > System Settings**.
2. In the **External Pools** card, click **Edit**.

The **External Service Pools** page appears.

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



The IP address must be associated with Nexus Dashboard Data Pool. A single persistent IP address is required to visualize and track EPs for a single site.

4. Configure SVI using FHRP for ND Data Interface and Underlay IP connectivity.

You can use **switch_freeform** policy on fabric Leaf 1.

Follow these steps to create a freeform policy.

- a. Choose **Manage > Fabrics**, then click on the required fabric.

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

- b. Click **Configuration Policies > Policies**.

- c. Click **Actions > Add policy**.

The **Create Policy** page appears.

- d. Choose the appropriate leaf switch from the **Switch List** drop-down list and click **Choose Template**.

- e. On **Select Policy Template** page, choose **switch_freeform** template and click **Select**.

Apply FHRP configurations and save the template.

Deploy the template configuration.

In this example, SVI 100 with HSRP gateway created on fabric Leaf 1. Similarly, repeat the steps for fabric Leaf 2.

The following is a configuration example:

```
feature hsrp
vlan 100
name EPL-Inband
interface Vlan100
  no shutdown
  no ip redirects
  ip address 192.168.100.252/24
  no ipv6 redirects
  ip router ospf 100 area 0.0.0.0
  hsrp 100
    ip 192.168.100.254
```

5. Verify IP reachability between Nexus Dashboard Data Interface and fabric switches.

6. Enable EPL at fabric level.

- a. In the fabric **Overview** page, choose **Actions > Configuration > Configure Endpoint Locator**.
- b. Choose the appropriate switches on the fabric hosting the Spine/Route Reflector RRs from the drop-down list.

Choose **Configure my Fabric** option for knob controls.

Whether BGP configuration will be pushed to the selected Spines/RRs as part of the enablement of the EPL feature. If the Spine/RR needs to be configured manually with a custom policy for the EPL BGP neighborship, then this option should be unchecked. For external fabrics that are only monitored and not configured on Nexus Dashboard this option is grayed out. As these fabrics are not configured on Nexus Dashboard.

Choose **Process MAC-Only Advertisements** option to enable processing of MAC-Only advertisements while configuring the EPL feature.



If EPL is enabled on a fabric with or without selecting the **Process Mac-Only Advertisements** checkbox and you want to toggle this selection later, you must first click **Database Clean-up** to delete endpoint data while EPL is still enabled. Attempting to clean the database while EPL is disabled will result in an error. After the cleanup is complete, you can disable and then re-enable EPL with the desired **Process Mac-Only Advertisements** setting.

Choose **Yes** in **Collect Additional Information** to enable collection of additional information such as PORT, VLAN, and VRF while enabling the EPL feature. To access additional information, NX-API must be supported and enabled on the switches, ToRs, and leafs. If you choose the **No** option, this information won't be collected and reported by EPL.



For all fabrics except external fabrics, NX-API is enabled by default. For external fabrics, you must enable NX-API in the external fabric settings, choose **Enable NX-API** checkbox in the **Advanced** tab of the External_Fabric_11_1 fabric template.


Click the **Preview** icon to view a template of the configuration that is pushed to the switches enabling EPL. This configuration can be copied and pasted on spines or border gateway devices to enable EPL on external monitored fabrics.

Once the appropriate selections are made and various inputs have been reviewed, click **Save Config** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message are displayed. Otherwise, EPL is successfully enabled. Once the EPL is enabled, the Persistent IP will be in-use.

Configuring Endpoint Locator for interconnected VXLAN fabrics

Before you begin:

The below figure enables EPL for interconnected VXLAN fabrics. The BGP peering's are established between the Spines/RRs of each interconnected VXLAN fabric and Nexus Dashboard EPL Container. The Persistent IPs are required based on the number of VXLAN EVPN sites. The Nexus Dashboard application hosted on Cisco ND Cluster is located on Site 1. The routing information to reach the Spines/RRs deployed in the remote site must be exchanged across the interconnected VXLAN fabrics. Once the BGP session is formed, local EPs of Fabric 2 can be visualized and tracked.



By default, Nexus Dashboard data Interface and Site 2 Spines/RRs loopback prefixes are not advertised across the BGWs. Therefore, prefixes must be exchanged using custom route maps and prefix lists across the sites. At the same time, route redistribution between OSPF and BGP is required as Spines/RRs loopback prefixes are part of OSPF protocol while BGWs peer with each other using BGP.

Follow these steps to configure endpoint locator for interconnected VXLAN fabrics.

1. You must configure persistent IP addresses on Cisco Nexus Dashboard. On Nexus Dashboard, choose **Admin > System Settings**.
2. In the **External Pools** card, click **Edit**.

The **External Service Pools** page appears.

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



Ensure that the IP addresses are associated with Nexus Dashboard Data Pool. Two persistent IP addresses are required to visualize and track EPs for a multisite with two member fabrics. One Persistent Data IP address is used as EPL container IP to establish BGP session with Site 1 fabric. A new Persistent IP address is configured that can be used to peer with Site 2 fabric.

4. Configure Route Redistribution for VXLAN EVPN Fabrics.

Route Redistribution for Fabric 1

The following `switch_freeform` policy can be used on Fabric 1 BGWs. To create a new **switch_freeform** policy, refer to the above examples.

The following example shows a sample configuration:

```

ip prefix-list site-2-rr seq 5 permit 20.2.0.1/32 >> Site 2 RR
ip prefix-list site-2-rr seq 6 permit 20.2.0.2/32 >> Site 2 RR
ip prefix-list epl-subnet seq 5 permit 192.168.100.0/24 >> EPL Subnet

route-map bgp-to-ospf permit 10
  match ip address prefix-list site-2-rr
route-map ospf-to-bgp permit 10
  match ip address prefix-list epl-subnet

router ospf 100
  redistribute bgp 100 route-map bgp-to-ospf

router bgp 100
  address-family ipv4 unicast
    redistribute ospf 100 route-map ospf-to-bgp

```

Route Redistribution for Fabric 2

The following `switch_freeform` policy can be used on Fabric 2 BGWs. To create a new **switch_freeform** policy, refer to the above examples.

The example below shows a sample configuration:

```

ip prefix-list site-2-rr seq 5 permit 20.2.0.1/32 >> Site 2 RR
ip prefix-list site-2-rr seq 6 permit 20.2.0.2/32 >> Site 2 RR
ip prefix-list epl-subnet seq 5 permit 192.168.100.0/24 >> EPL Subnet

route-map bgp-to-ospf permit 10
  match ip address prefix-list epl-subnet
route-map ospf-to-bgp permit 10
  match ip address prefix-list site-2-rr

router ospf 200
  redistribute bgp 200 route-map bgp-to-ospf

router bgp 200
  address-family ipv4 unicast
    redistribute ospf 200 route-map ospf-to-bgp

```

5. To configure EPL, in the fabric **Overview** page, choose **Actions > Configuration > Configure Endpoint Locator**.
6. Choose the appropriate switches on the fabric hosting the Spine/Route Reflector RRs from the drop-down list.

Once the appropriate selections are made and various inputs have been reviewed, click **Save**

Config to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled. Once the EPL is enabled, the Persistent IP will be in-use.

You can view EPL enabled for fabric-1 and fabric-2 successfully. To view and track EPs, see [Monitoring Endpoint Locator](#).

Configuring Endpoint Locator for vPC Fabric Peering Switches

Networks Administrator can create vPC between a pair of switches using a Physical Peer Link or Virtual Peer link. vPC Fabric Peering provides an enhanced dual-homing access solution without the overhead of wasting physical ports for vPC Peer Link. For Virtual Peer link, EPL can still be connected to vPC pair of Leaf switches for the link and node-level redundancy. However, VXLAN VLAN (Anycast Gateway) as the First hop for EPL will be used. The loopback0 address of Spines/RRs is reachable only via VXLAN Underlay, while VXLAN VLAN will be part of a Tenant VRF. Therefore, to establish IP communication, route-leaking is configured between Tenant VRF and Default VRF. For more information, refer to vPC Fabric Peering section.

Follow these steps to configure endpoint locator for vPC Fabric Peering switches.

1. You must configure persistent IP addresses on Cisco Nexus Dashboard. On Nexus Dashboard, choose **Admin > System Settings**.
2. In the **External Pools** card, click **Edit**.

The **External Service Pools** page appears.

3. In the **Persistent data IPs** area, click **+ Add IP Address** and enter the necessary persistent IP addresses. Click the checkmark for each IP address that you enter.
4. Create a Tenant VRF and Anycast Gateway on the vPC fabric peering switches.
5. Configure Route-leaking between Tenant VRF and Default VRF.

Advertise from Tenant VRF to Default VRF.

The following switch_freeform policy can be used on fabric Leaf where ND is connected.

```
ip prefix-list vrf-to-default seq 5 permit 192.168.100.0/24 >> EPL subnet
route-map vrf-to-default permit 10
  match ip address prefix-list vrf-to-default
vrf context epl_inband
  address-family ipv4 unicast
    export vrf default map vrf-to-default allow-vpn
router ospf UNDERLAY
  redistribute bgp 200 route-map vrf-to-default
```

NOTE: For the hosts behind the orphan ports, the prefix-list used for route leaking in the route-map must include 'ge' keyword at the end for longer prefix matches, such as the /32 of a host route.

For example:

```
ip prefix-list vrf-to-default seq 5 permit 192.168.100.0/24 ge 24
```

Advertise from Default VRF to Tenant VRF.

The following switch_freeform policy can be used on fabric Leaf where ND is connected.

```
ip prefix-list default-to-vrf seq 5 permit 20.2.0.3/32 >> Spine loopback IP
ip prefix-list default-to-vrf seq 6 permit 20.2.0.4/32 >> Spine loopback IP
route-map default-to-vrf permit 10
  match ip address prefix-list default-to-vrf
vrf context epl_inband
  address-family ipv4 unicast
    import vrf default map default-to-vrf
    router bgp 200
  address-family ipv4 unicast
    redistribute ospf UNDERLAY route-map default-to-vrf
```

6. Enable EPL at fabric level.

- a. To configure EPL, in the fabric **Overview** page, choose **Actions > Configuration > Configure Endpoint Locator**.
- b. Choose the appropriate switches on the fabric hosting the Spine/Route Reflector RRs from the drop-down list.

Once the appropriate selections are made and various inputs have been reviewed, click **Save Config** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled. Once the EPL is enabled, the Persistent IP will be in-use.

Configuring Endpoint Locator for External Fabrics

In addition to Easy fabrics, Nexus Dashboard allows you to enable EPL for VXLAN EVPN fabrics comprising of switches that are imported into the external fabric. The external fabric can be in managed mode or monitored mode, based on the selection of **Fabric Monitor Mode** flag in the **External Fabric** Settings. For external fabrics that are only monitored and not configured by Nexus Dashboard, this flag is disabled. Therefore, you must configure BGP sessions on the Spine(s) via OOB or using the CLI. To check the sample template, click  icon to view the configurations required while enabling EPL.

In case the **Fabric Monitor Mode** checkbox in the External Fabric settings is unchecked, then EPL can still configure the spines/RRs with the default **Configure my fabric** option. However, disabling EPL would wipe out the router bgp config block on the spines/RRs. To prevent this, the BGP policies must be manually created and pushed onto the selected spines/RRs.

Configuring Endpoint Locator for eBGP EVPN Fabrics

You can enable EPL for VXLAN EVPN fabrics, where eBGP is employed as the underlay routing protocol. Note that with an eBGP EVPN fabric deployment, there is no traditional RR similar to iBGP.

The reachability of the in-band subnet must be advertised to the spines that behave as Route Servers.

To configure EPL for eBGP EVPN fabrics from the Cisco Nexus Dashboard Web UI, perform the following steps:

1. Choose **Manage > Fabrics**.

Select the fabric to configure eBGP on or create an eBGP fabric with the **VXLAN** template.

2. Use the **leaf_bgp_asn** policy to configure unique ASNs on all leaf switches.
3. Add the **ebgp_overlay_leaf_all_neighbor** policy to each leaf.

Fill **Spine IP List** with the spines' BGP interface IP addresses, typically the loopback0 IP addresses.

Fill **BGP Update-Source Interface** with the leaf's BGP interface, typically loopback0.

4. Add the **ebgp_overlay_spine_all_neighbor** policy to each spine.

Fill **Leaf IP List** with the leaves' BGP interface IPs, typically the loopback0 IPs.

Fill **Leaf BGP ASN** with the leaves' ASNs in the same order as in **Leaf IP List**.

Fill **BGP Update-Source Interface** with the spine's BGP interface, typically loopback0.

After the in-band connectivity is established, the enablement of the EPL feature remains identical to what is listed so far. EPL becomes a iBGP neighbor to the Route Servers running on the spines.

Flushing the Endpoint Database

After you enable the Endpoint Locator feature, you can clean up or flush all the Endpoint information. This allows starting from a clean-slate with respect to ensuring no stale information about any endpoint is present in the database. After the database is clean, the BGP client re-populates all the endpoint information learnt from the BGP RR. You can flush the endpoint database even if you have not re-enabled the EPL feature on a fabric on which the EPL feature was previously disabled.

To flush all the Endpoint Locator information from the Cisco Nexus Dashboard Web UI, perform the following steps:

1. In the fabric **Overview** page, choose **Actions > Configuration > Configure Endpoint Locator**
2. Click **Database Cleanup**.
3. Click **Delete** to continue or **Cancel** to abort.

Monitoring Endpoint Locator

Information about the Endpoint Locator (EPL) is displayed on a single landing page or dashboard. The dashboard displays an almost real-time view of data, refreshed every 30 seconds, showing all the active endpoints on a single pane. From this dashboard, you can access Endpoint History, Endpoint Search, and Endpoint Life.



The EPL snapshot summary counts for endpoints, VRFs, and networks are updated every hour on a physical Cisco Nexus Dashboard cluster and every six hours on a virtual Cisco Nexus Dashboard cluster. Snapshot comparisons on virtual clusters reflect a six-hour granularity instead of a one-hour granularity.

Disabling Endpoint Locator

Follow these steps to disable endpoint locator from the Cisco Nexus Dashboard UI.

1. In the fabric **Overview** page, choose **Actions > Configuration > Configure Endpoint Locator**.

The **Endpoint Locator** page appears.

2. Click **Disable**.

When EPL is disabled, BGP neighbor configuration is removed from the switches and the EPL worker pod is terminated. Historical endpoint data is retained in the databases unless explicitly cleared using the Clear Fabric Data option.

Copyright

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

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

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

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

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

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

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

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

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

Americas Headquarters

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

<https://www.cisco.com>

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883