



Management


This chapter contains the following sections:

- [Management Workflows, on page 1](#)
- [Adding Management Access, on page 2](#)
- [Exporting Tech Support, Statistics, and Core Files, on page 11](#)
- [Overview, on page 13](#)
- [Backing up, Restoring, and Rolling Back Controller Configuration, on page 19](#)
- [Using the Cisco APIC Troubleshooting Tools, on page 29](#)

Management Workflows

ACI Management Access Workflows

This workflow provides an overview of the steps required to configure management connectivity to switches in the ACI fabric.



1. Prerequisites

- Ensure that you have read/write access privileges to the infra security domain.
- Ensure that the target leaf switches with the necessary interfaces are available.

2. Configure the ACI Leaf Switch Access Ports

Choose which of these management access scenarios you will use:

- For **in-band** management, follow the suggested topics for in-band configuration in the *APIC Basic Configuration Guide*.
- For **out-of-band** management, follow the suggested topics for out-of-band configuration in the *APIC Basic Configuration Guide*.

Suggested topics

For additional information, see the following topics in the *APIC Basic Configuration Guide*:

- Configuring In-Band Management Access Using the Advanced GUI
- Configuring In-Band Management Access Using the NX-OS Style CLI
- Configuring In-Band Management Access Using the REST API
- Configuring Out-of-Band Management Access Using the Advanced GUI
- Configuring Out-of-Band Management Access Using the NX-OS Style CLI
- Configuring Out-of-Band Management Access Using the REST API

Adding Management Access

Configuring the external management instance profile under the management tenant for in-band has no effect on the protocols that are configured under the fabric-wide communication policies. The subnets and contracts specified under the external management instance profile do not affect HTTP/HTTPS or SSH/Telnet. Beginning with the 6.0(2) release, Telnet is not supported.

Adding Management Access in the GUI

A Cisco Application Policy Infrastructure Controller (APIC) has two routes to reach the management network: one is by using the in-band management interface and the other is by using the out-of-band management interface.

The in-band management network allows Cisco APIC to communicate with the leaf switches and with the outside using the Cisco Application Centric Infrastructure (ACI) fabric, and it makes it possible for external management devices to communicate with the Cisco APIC or the leaf switches and spine switches using the fabric itself.

The out-of-band management network configuration defines the configuration of the management port on the controllers, the leaf switches and the spine switches.

The Cisco APIC controller always selects the in-band management interface over the out-of-band management interface, if the in-band management interface is configured. The out-of-band management interface is used only when the in-band management interface is not configured or if the destination address is on the same subnet as the out-of-band management subnet of the Cisco APIC.

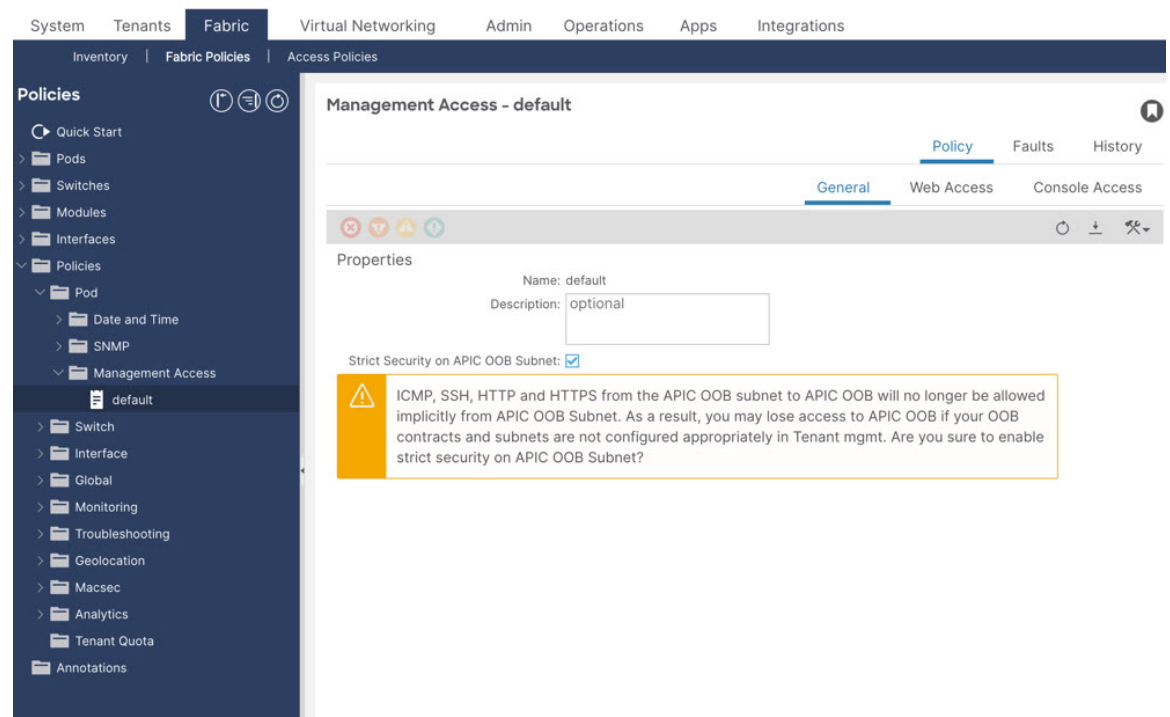
Cisco ACI has the ability to program routes for in-band management based on the subnet configuration on the bridge domains in the management tenant and in-band VRF instance. These routes will be deleted when the subnet configuration is deleted from the bridge domains.

The Cisco APIC out-of-band management connection link must be 1 Gbps.

Access to APIC using the OOB management IP can be restricted by configuring contracts and allowed subnets. However, IP addresses which belong to the APIC OOB subnet are always allowed to access APIC OOB using ICMP, SSH, HTTP, HTTPS, and TCP 4200 regardless of the configuration of contracts and subnets. This safeguard is in place to prevent accidental loss of APIC access due to misconfiguration. From Cisco APIC Release 6.1(1), you can remove this safeguard by enabling the **Strict Security on APIC OOB Subnet** option while configuring the default Management Access, which is located at **Fabric > Fabric Policies > Policies > Pod > Management Access** on the APIC GUI.

As soon as you enable strict mode, a warning is displayed on the Management Access screen (as shown below). A similar warning about enabling strict mode is displayed on the **External Management Network Profile** screen too (located under **Tenants > mgmt > External Management Network Instance Profiles**).

Figure 1: Warning displayed after enabling Strict Security mode



If you enable **Strict Security on APIC OOB Subnet** with a contract and a subnet (in the External Management Network Instance Profile) that does not allow your IP address to access the APIC, you will lose access to the APIC through the OOB IP address. You will be unable to access the APIC from the same network as the

APIC OOB subnet either. In such a case, login to the console and use the `acidiag enableoobrecovery` command to disable **Strict Security on APIC OOB Subnet**.



Note Duplicate IP addresses and firewalls that cache ARP information are not supported on the management network. The presence of these conditions can result in the complete loss of Cisco APIC management access following an upgrade.

IPv4/IPv6 Addresses and In-Band Policies

In-band management addresses can be provisioned on the APIC controller only through a policy (Postman REST API, NX-OS Style CLI, or GUI). Additionally, the in-band management addresses must be configured statically on each node.

IPv4/IPv6 Addresses in Out-of-Band Policies

Out-of-band management addresses can be provisioned on the APIC controller either at the time of bootstrap or by using a policy (Postman REST API, NX-OS Style CLI, GUI). Additionally, the out-of-band management addresses must be configured statically on each node or by specifying a range of addresses (IPv4/IPv6) to the entire cluster. IP addresses are randomly assigned from a range to the nodes in the cluster.

IPv6 Table Modifications to Mirror the Existing IP Tables Functionality

All IPv6 tables mirror the existing IP tables functionality, except for Network Address Translation (NAT).

Existing IP Tables

1. Earlier, every rule in the IPv6 tables were executed one at a time and a system call was made for every rule addition or deletion.
2. Whenever a new policy was added, rules were appended to the existing IP tables file and no extra modifications were done to the file.
3. When a new source port was configured in the out-of-band policy, it added source and destination rules with the same port number.

Modifications to IP Tables

1. When IP tables are created, they are first written into hash maps that are then written into intermediate file `IP tables-new` which are restored. When saved, a new IP tables file is created in the `/etc/sysconfig/` folder. You can find both these files at the same location. Instead of making a system call for every rule, you must make a system call only while restoring and saving the file.
2. When a new policy is added instead of appending it to the file, an IP table is created from scratch, that is by loading default policies into the hashmaps, checking for new policies, and adding them to hashmaps. Later, they are written to the intermediate file (`/etc/sysconfig/iptables-new`) and saved.
3. It is not possible to configure source ports alone for a rule in out-of-band policy. Either destination port or source port along with a destination port can be added to the rules.

- When a new policy is added, a new rule will be added to the IP tables file. This rule changes the access flow of IP tables default rules.

```
-A INPUT -s <OOB Address Ipv4/Ipv6> -j apic-default
```

- When a new rule is added, it presents in the IP tables-new file and not in the IP tables file, and it signifies that there is some error in the IP tables-new file. Only if the restoration is successful, the file is saved and new rules are seen in the IP tables file.



Note

- If only IPv4 is enabled, do not configure an IPv6 policy.
 - If only IPv6 is enabled, do not configure an IPv4 policy.
 - If both IPv4 and IPv6 are enabled and a policy is added, it will be configured to both the versions . So when you add an IPv4 subnet, it will be added to IP tables and similarly an IPv6 subnet is added to IPv6 tables.
-

Management Access Guidelines and Restrictions

- vzAny is supported as a consumer of a shared service but is not supported as a provider of a shared service. vzAny shared service consumer and vzAny provider is not supported.
- When configuring out-of-band management access, logging options for an out-of-band contract (enabling and viewing ACL contract and permit/deny logs) is not supported.
- An in-band management address must be configured for a leaf node in order to push the in-band management VRF to a leaf node.
- A bridge domain subnet IP address in the in-band management VRF can be assigned as a secondary IP address unless "Make this IP address primary" is selected for a gateway subnet.
- The following ports cannot be denied in an out-of-band contract:
 - proto icmp, rate-limited, not configurable
 - tcp dpt: 22, rate-limited, not configurable
 - tcp dpt: 80, by default no listening process
 - tcp dpt: 443, default UI/API
 - tcp dpt: 4200, SSH access through web, by default no listening process

When you define a subnet under the external network instance profiles, the port list above is restricted to the sources in the configured OOB subnet.

The OOB contract does not take effect for the corresponding address family if the IPv4 or IPv6 subnet is not defined under the external network instance profiles.

To enable the OOB contract for both IPv4 and IPv6, you must configure at least one IPv4 and one IPv6 subnet under the external network instance profiles.

For SNMP on leaf switches and spine switches, the **Client Entries** subnets that are configured under **Fabric Policies > Pod > SNMP** are matched before the OOB contract. If there are no subnets configured under **Client Entries**, any source is allowed for SNMP. For example, UDP **dpt:161**.

By default, no IP address is allocated to the management interface of leaf switches and spine switches. However, once an IP address is assigned, there are some ports that cannot be denied in an out-of-band contract. They are required for the built-in features of ACI. For example, NTP, DHCP, ICMP, and so on.

Subnets that are defined under the external network instance profiles are applicable only to APIC. On leaf switches and spine switches, any source (0.0.0.0/0) is allowed.

- A spine switch does not resolve ARP on the in-band management IP address. Due to this, any device in the in-band management network cannot communicate with the spine switch. Access to a spine switch is only possible over a Layer 3 network.

-

Configuring In-Band and Out-of-Band Management Access with Wizards

In APIC, release 3.1(x), wizards were added to simplify configuring management access. You can still use the other methods of configuring management access included in this document.

Procedure

-
- Step 1** To configure **In-Band Management Access**, perform the following steps:
- a) On the menu bar, click **Tenants > mgmt**.
 - b) Expand **Quick Start**.
 - c) Click **In-Band Management Access > Configure In-Band Management Access > Start**.
 - d) Follow the instructions to add the **Nodes** in the management network, the **IP addresses** for the nodes, communication filters for the **Connected Devices**, and communication filters for **Remote Attached Devices**.
- Step 2** To configure **Out-of-Band Management Access**, perform the following steps:
- a) On the menu bar, click **Tenants > mgmt**.
 - b) Expand **Quick Start**.
 - c) Click **Out-of-Band Management Access > Configure Out-of-Band Management Access > Start**.
 - d) Follow the instructions to add the **Nodes** in the out-of-band management network, the **IP addresses** for the nodes, subnets allowed for the **External Hosts**, and communication filters that will determine communication for **Access**.
-

Configuring In-Band Management Access Using the Cisco APIC GUI



Note IPv4 and IPv6 addresses are supported for in-band management access. IPv6 configurations are supported using static configurations (for both in-band and out-of-band). IPv4 and IPv6 dual in-band and out-of-band configurations are supported only through static configuration. For more information, see the KB article, *Configuring Static Management Access in Cisco APIC*.

Procedure

- Step 1** On the menu bar, choose **Fabric > Access Policies**.
- Step 2** In the **Navigation** pane, right-click **Interfaces** and choose **Configure Interface, PC and VPC**.
- Step 3** In the **Configure Interface, PC, and VPC** dialog box, to configure switch ports connected to Cisco Application Policy Infrastructure Controllers (APICs), perform the following actions:
- Click the large + icon next to the switch diagram to create a new profile and configure VLANs for the Cisco APIC.
 - From the **Switches** field drop-down list, check the check boxes for the switches to which the Cisco APICs are connected. (leaf1 and leaf2).
 - In the **Switch Profile Name** field, enter a name for the profile (apicConnectedLeaves).
 - Click the + icon to configure the ports.
 - Verify that in the **Interface Type** area, the **Individual** radio button is selected.
 - In the **Interfaces** field, enter the ports to which the Cisco APICs are connected.
 - In the **Interface Selector Name** field, enter the name of the port profile (apicConnectedPorts).
 - In the **Interface Policy Group** field, click the **Create One** radio button.
 - In the **Attached Device Type** field, choose the appropriate device type to configure the domain (Bare Metal).
 - In the **Domain** field, click the **Create One** radio button.
 - In the **Domain Name** field, enter the domain name. (**inband**)
 - In the **VLAN** field, choose the **Create One** radio button.
 - In the **VLAN Range** field, enter the VLAN range. Click **Save**, and click **Save** again. Click **Submit**.
- Step 4** In the **Navigation** pane, right-click **Switch Policies** and choose **Configure Interface, PC and VPC**.
- Step 5** In the **Configure Interface, PC, and VPC** dialog box, perform the following actions:
- Click the large + icon next to the switch diagram to create a new profile and configure VLANs for the server.
 - In the **Switches** field, from drop-down list, check the check boxes for the switches to which the servers are connected. (leaf1).
 - In the **Switch Profile Name** field, enter a name for the profile (vmmConnectedLeaves).
 - Click the + icon to configure the ports.
 - Verify that in the **Interface Type** area, the **Individual** radio button is selected.
 - In the **Interfaces** field, enter the ports to which the servers are connected (1/40).
 - In the **Interface Selector Name** field, enter the name of the port profile.
 - In the **Interface Policy Group** field, click the **Create One** radio button.

- i) In the **Attached Device Type** field, choose the appropriate device type to configure the domain (Bare Metal).
- j) In the **Domain** field, from the drop-down list click the **Choose One** radio button
- k) From the **Physical Domain** drop-down list, choose the domain created earlier.
- l) In the **Domain Name** field, enter the domain name.
- m) Click **Save**, and click **Save** again.

Step 6 In the **Configure Interface, PC, and VPC** dialog box, click **Submit**.

Step 7 On the menu bar, click **TENANTS > mgmt**. In the **Navigation** pane, expand **Tenant mgmt > Networking > Bridge Domains** to configure the bridge domain on the in-band connection.

Step 8 Expand the in-band bridge domain (inb). Right-click **Subnets**. Click **Create Subnet** and perform the following actions to configure the in-band gateway:

- a) In the **Create Subnet** dialog box, in the **Gateway IP** field, enter the in-band management gateway IP address and mask.
- b) Click **Submit**.

Step 9 In the **Navigation** pane, expand **Tenant mgmt > Node Management EPGs**. Right-click **Node Management EPGs** and choose **Create In-Band Management EPG**. Perform the following actions to set the VLAN on the in-band EPG used to communicate with the Cisco APIC:

- a) In the **Name** field, enter the in-band management EPG name.
- b) In the **Encap** field, enter the VLAN (vlan-10).
- c) From the **Bridge Domain** drop-down field, choose the bridge domain. Click **Submit**.
- d) In the **Navigation** pane, choose the newly created in-band EPG.
- e) Expand **Provided Contracts**. In the **Name** field, from the drop-down list, choose the default contract to enable EPG to provide the default contract that will be consumed by the EPGs on which the VMM servers are located.
- f) Click **Update**, and click **Submit**.

Step 10 In the **Navigation** pane, right-click **Node Management Addresses** and click **Create Node Management Addresses**, and perform the following actions to configure the IP addresses to be assigned to the Cisco APICs in the fabric:

- a) In the **Create Node Management Addresses** dialog box, in the **Policy Name** field, enter the policy name (apicInb).
- b) In the **Nodes** field, **Select** column, check the check boxes for the nodes that will be part of this fabric (apic1, apic2, apic3).
- c) In the **Config** field, check the **In-Band Addresses** check box.
- d) In the **Node Range** fields, enter the range.
- e) In the **In-Band IP Addresses** area, in the **In-Band Management EPG** field, from the drop-down list, choose default. This associates the default in-band Management EPG.
- f) In the **In-Band IP Addresses** and **Gateway** fields, enter the IPv4 or IPv6 addresses as desired.
- g) Click **Submit**. The IP addresses for the Cisco APICs are now configured.

Step 11 In the **Navigation** pane, right-click **Node Management Addresses**. Click **Create Node Management Addresses**, and perform the following actions to configure the IP addresses for the leaf and spine switches in the fabric:

- a) In the **Create Node Management Addresses** dialog box, in the **Policy Name** field, enter the policy name (switchInb).
- b) In the **Nodes** field, **Select** column, check the check boxes next to the nodes that will be part of this fabric (leaf1, leaf2, spine1, spine2).

- c) In the **Config** field, click the **In-Band Addresses** checkbox.
- d) In the **Node Range** fields, enter the range.
- e) In the **In-Band IP Addresses** area, in the **In-Band Management EPG** field, from the drop-down list, choose default. The default in-band management EPG is now associated.
- f) In the **In-Band IP Addresses** and **Gateway** fields, enter the IPv4 or IPv6 addresses as desired.
- g) Click **Submit**. In the **Confirm** dialog box, click **Yes**. The IP addresses for the leaf and spine switches are now configured.

Step 12 In the **Navigation** pane, under **Node Management Addresses**, click the Cisco APIC policy name (apicInb) to verify the configurations. In the **Work** pane, the IP addresses assigned to various nodes are displayed.

Step 13 In the **Navigation** pane, under **Node Management Addresses**, click the switches policy name (switchInb). In the **Work** pane, the IP addresses that are assigned to switches and the gateway addresses they are using are displayed.

Note You can make out-of-band management access the default management connectivity mode for the Cisco APIC server by clicking **System > System Settings > APIC Connectivity Preferences**. Then on the **Connectivity Preferences** page, click **inband**.

Configuring Out-of-Band Management Access Using the Cisco APIC GUI



Note IPv4 and IPv6 addresses are supported for out-of-band management access.

You must configure out-of-band management access addresses for the leaf and spine switches as well as for the Cisco APIC.

Before you begin

The Cisco Application Policy Infrastructure Controller (APIC) out-of-band management connection link must be 1 Gbps.

Procedure

Step 1 On the menu bar, choose **Tenants > mgmt**. In the **Navigation** pane, expand **Tenant mgmt**.

Step 2 Right-click **Node Management Addresses**, and click **Create Node Management Addresses**.

Step 3 In the **Create Node Management Addresses** dialog box, perform the following actions:

- a) In the **Policy Name** field, enter a policy name (switchOob).
- b) In the **Nodes** field, check the check boxes next to the appropriate leaf and spine switches (leaf1, leaf2, spine1).
- c) In the **Config** field, check the check box for **Out of-Band Addresses**.

Note The **Out-of-Band IP addresses** area is displayed.

- d) In the **Out-of-Band Management EPG** field, choose the EPG from the drop-down list (default).

- e) In the **Out-Of-Band Gateway** field, enter the IP address and network mask for the external out-of-band management network.
- f) In the **Out-of-Band IP Addresses** field, enter the range of desired IPv4 or IPv6 addresses that will be assigned to the switches. Click **Submit**.

The node management IP addresses are configured.

Step 4 In the **Navigation** pane, expand **Node Management Addresses**, and click the policy that you created. In the **Work** pane, the out-of-band management addresses are displayed against the switches.

Step 5 In the **Navigation** pane, expand **Contracts > Out-of-Band Contracts**.

Step 6 Right-click **Out-of-Band Contracts**, and click **Create Out-of-Band Contract**.

Step 7 In the **Create Out-of-Band Contract** dialog box, perform the following tasks:

- a) In the **Name** field, enter a name for the contract (oob-default).
- b) Expand **Subjects**. In the **Create Contract Subject** dialog box, in the **Name** field, enter a subject name (oob-default).
- c) Expand **Filter Chain**, and in the **Name** field, from the drop-down list, choose the name of the filter (default). Click **Update**, and click **OK**.
- d) In the **Create Out-of-Band Contract** dialog box, click **Submit**.

An out-of-band contract that can be applied to the out-of-band EPG is created.

Step 8 In the **Navigation** pane, expand **Node Management EPGs > Out-of-Band EPG - default**.

Step 9 In the **Work** pane, expand **Provided Out-of-Band Contracts**.

Step 10 In the **OOB Contract** column, from the drop-down list, choose the out-of-band contract that you created (oob-default). Click **Update**, and click **Submit**.

The contract is associated with the node management EPG.

Step 11 In the **Navigation** pane, right-click **External EPG**, and click **Create External Management Entity Instance**.

Step 12 In the **Create External Management Entity Instance** dialog box, perform the following actions:

- a) In the **Name** field, enter a name (oob-mgmt-ext).
- b) Expand the **Consumed Out-of-Band Contracts** field. From the **Out-of-Band Contract** drop-down list, choose the contract that you created (oob-default). Click **Update**.

Choose the same contract that was provided by the out-of-band management.

- c) In the **Subnets** field, enter the subnet address. Click **Submit**.

Only the subnet addresses you choose here will be used to manage the switches. The subnet addresses that are not included cannot be used to manage the switches.

The node management EPG is attached to the external EPG. The out-of-band management connectivity is configured.

Note You can make out-of-band management access the default management connectivity mode for the Cisco APIC server by clicking **System > System Settings > APIC Connectivity Preferences**. Then on the **Connectivity Preferences** page, click **ooband**.

Exporting Tech Support, Statistics, and Core Files

About Exporting Files

An administrator can configure export policies in the APIC to export statistics, technical support collections, faults and events, to process core files and debug data from the fabric (the APIC as well as the switch) to any external host. The exports can be in a variety of formats, including XML, JSON, web sockets, secure copy protocol (SCP), or HTTP. You can subscribe to exports in streaming, periodic, or on-demand formats.

An administrator can configure policy details such as the transfer protocol, compression algorithm, and frequency of transfer. Policies can be configured by users who are authenticated using AAA. A security mechanism for the actual transfer is based on a username and password. Internally, a policy element handles the triggering of data.

File Export Guidelines and Restrictions

- HTTP export and the streaming API format is supported only with statistics information. Core and tech support data are not supported.
- The destination IP address for exported files cannot be an IPv6 address.
- Do not trigger tech support from more than five nodes simultaneously, especially if they are to be exported into the Cisco Application Policy Infrastructure Controller (APIC) or to an external server with insufficient bandwidth and compute resources.
- To collect tech support from all of the nodes in the fabric periodically, you must create multiple policies. Each policy must cover a subset of the nodes and should be scheduled to trigger in a staggered way (at least 30 minutes apart).
- Do not schedule more than one tech support policy for the same node on the Cisco APIC. Running multiple instances of tech support policies on the same node at the same time can result in a huge consumption of Cisco APIC or switch CPU cycles and the other resources.
- We recommend that you use the regular tech support policy for the nodes placed in maintenance mode instead of the on-demand tech support policy.
- The status of an on-going tech support for the nodes in maintenance mode will not be available in the Cisco APIC GUI in the **Admin > Tech Support > *policy_name* > Operational > Status** section. Based on your selection of **Export to Controller** or **Export Destination** in the tech support policy, you can verify the controller (`/data/techsupport`) or the destination server to confirm that the tech support is being captured.
- Tech support collection from the Cisco APIC can time out when the cores on a leaf switch are busy. The cores can become busy if routing processes such as BGP and platform processes such as HAL hog the CPU. If the tech support collection times out, check for the CPU utilization to see if there is a CPU hog. If there is, you can collect the tech support on the leaf switch directly to avoid the timeout issues.

Creating a Remote Location for Exporting Files

This procedure configures the host information and file transfer settings for a remote host that will receive exported files.

Procedure

- Step 1** In the menu bar, click **Admin**.
- Step 2** In the submenu bar, click **Import/Export**.
- Step 3** In the **Navigation** pane, expand **Export Policies**.
- Step 4** Right-click **Remote Locations** and choose **Create Remote Path of a File**.
- Step 5** In the **Create Remote Path of a File** dialog box, perform the following actions:
- In the **Name** field, enter a name for the remote location.
 - In the **Host Name/IP** field, enter an IP address or a fully qualified domain name for the destination host.
 - In the **Protocol** field, click the radio button for the desired file transfer protocol.
 - In the **Remote Path** field, type the path where the file will be stored on the remote host.
 - Enter a username and password for logging in to the remote host and confirm the **Password**.
 - From the **Management EPG** drop-down list, choose the management EPG.
 - Click **Submit**.
-

Sending an On-Demand Tech Support File Using the GUI

Procedure

- Step 1** In the menu bar, click **Admin**.
- Step 2** In the submenu bar, click **Import/Export**.
- Step 3** In the **Navigation** pane, expand **Export Policies**.
- Step 4** Right-click **On-demand Tech Support** and choose **Create On-demand Tech Support**.
The **Create On-demand Tech Support** dialog box appears.
- Step 5** Enter the appropriate values in the fields of the **Create On-demand Tech Support** dialog box.
- Note** For an explanation of a field, click the help icon in the **Create On-demand Tech Support** dialog box. The help file opens to a properties description page.
- Step 6** Click **Submit** to send the tech support file.
- Note** On-demand tech support files can be saved to another APIC to balance storage and CPU requirements. To verify the location, click on the On-demand Tech Support policy in the **Navigation** pane, then click the **OPERATIONAL** tab in the **Work** pane. The controller is displayed in the **EXPORT LOCATION** field.

Step 7 Right-click the policy name and choose **Collect Tech Support**.

Step 8 Choose **Yes** to begin collecting tech support information.

Overview

This topic provides information on:

- How to use configuration Import and Export to recover configuration states to the last known good state using the Cisco APIC
- How to encrypt secure properties of Cisco APIC configuration files

You can do both scheduled and on-demand backups of user configuration. Recovering configuration states (also known as "roll-back") allows you to go back to a known state that was good before. The option for that is called an Atomic Replace. The configuration import policy (configImportP) supports atomic + replace (importMode=atomic, importType=replace). When set to these values, the imported configuration overwrites the existing configuration, and any existing configuration that is not present in the imported file is deleted. As long as you do periodic configuration backups and exports, or explicitly trigger export with a known good configuration, then you can later restore back to this configuration using the following procedures for the CLI, REST API, and GUI.

For more detailed conceptual information about recovering configuration states using the Cisco APIC, please refer to the *Cisco Application Centric Infrastructure Fundamentals Guide*.

The following section provides conceptual information about encrypting secure properties of configuration files:

Configuration File Encryption

As of release 1.1(2), the secure properties of APIC configuration files can be encrypted by enabling AES-256 encryption. AES encryption is a global configuration option; all secure properties conform to the AES configuration setting. It is not possible to export a subset of the ACI fabric configuration such as a tenant configuration with AES encryption while not encrypting the remainder of the fabric configuration. See the *Cisco Application Centric Infrastructure Fundamentals*, "Secure Properties" chapter for the list of secure properties.

The APIC uses a 16 to 32 character passphrase to generate the AES-256 keys. The APIC GUI displays a hash of the AES passphrase. This hash can be used to see if the same passphrases was used on two ACI fabrics. This hash can be copied to a client computer where it can be compared to the passphrase hash of another ACI fabric to see if they were generated with the same passphrase. The hash cannot be used to reconstruct the original passphrase or the AES-256 keys.

Observe the following guidelines when working with encrypted configuration files:

- Backward compatibility is supported for importing old ACI configurations into ACI fabrics that use the AES encryption configuration option.



Note Reverse compatibility is not supported; configurations exported from ACI fabrics that have enabled AES encryption cannot be imported into older versions of the APIC software.

- Always enable AES encryption when performing fabric backup configuration exports. Doing so will assure that all the secure properties of the configuration will be successfully imported when restoring the fabric.



Note If a fabric backup configuration is exported without AES encryption enabled, none of the secure properties will be included in the export. Since such an unencrypted backup would not include any of the secure properties, it is possible that importing such a file to restore a system could result in the administrator along with all users of the fabric being locked out of the system.

- The AES passphrase that generates the encryption keys cannot be recovered or read by an ACI administrator or any other user. The AES passphrase is not stored. The APIC uses the AES passphrase to generate the AES keys, then discards the passphrase. The AES keys are not exported. The AES keys cannot be recovered since they are not exported and cannot be retrieved via the REST API.
- The same AES-256 passphrase always generates the same AES-256 keys. Configuration export files can be imported into other ACI fabrics that use the same AES passphrase.
- For troubleshooting purposes, export a configuration file that does not contain the encrypted data of the secure properties. Temporarily turning off encryption before performing the configuration export removes the values of all secure properties from the exported configuration. To import such a configuration file that has all secure properties removed, use the import merge mode; do not use the import replace mode. Using the import merge mode will preserve the existing secure properties in the ACI fabric.
- By default, the APIC rejects configuration imports of files that contain fields that cannot be decrypted. Use caution when turning off this setting. Performing a configuration import inappropriately when this default setting is turned off could result in all the passwords of the ACI fabric to be removed upon the import of a configuration file that does not match the AES encryption settings of the fabric.



Note Failure to observe this guideline could result in all users, including fabric administrations, being locked out of the system.

Configuring a Remote Location Using the GUI

This procedure explains how to create a remote location using the APIC GUI.

Procedure

-
- Step 1** On the menu bar, choose **ADMIN > Import/Export**.

- Step 2** In the navigation pane, right-click **Remote Locations** and choose **Create Remote Location**. The **Create Remote Location** dialog appears.
- Step 3** Enter the appropriate values in the **Create Remote Location** dialog fields.
- Note** For an explanation of a field, click the 'i' icon to display the help file.
- Step 4** When finished entering values in the **Create Remote Location** dialog fields, click **Submit**. You have now created a remote location for backing up your data.
-

Configuring an Export Policy Using the GUI

This procedure explains how to configure an export policy using the Cisco Application Policy Infrastructure Controller (APIC) GUI. Use the following procedure to trigger a backup of your data.



Note The **Maximum Concurrent Nodes** value that is configured in a scheduler policy determines the number of configuration export policies to act at the time that is specified in the scheduler policy.

For example, if the **Maximum Concurrent Nodes** is set to **1** in a scheduler policy and you have configured two export policies, both utilizing the same scheduler policy, one export policy is successful and other fails. However, if the **Maximum Concurrent Nodes** is set to **2**, both configurations are successful.

When the user is logged in with read-only privileges, Tech Support data can still be exported by right-clicking on the On-Demand Tech Support or Configuration Export polices and choosing **Trigger**.

Procedure

- Step 1** On the menu bar, choose **Admin > Import/Export**.
- Step 2** In the **Navigation** pane, right-click **Export Policies** and choose **Create Configuration Export Policy**. The **Create Configuration Export Policy** dialog appears.
- Step 3** Enter the appropriate values in the **Create Configuration Export Policy** dialog fields.
- For an explanation of a field, click the help (?) icon to display the help file.
- Step 4** After you finish entering values in the **Create Configuration Export Policy** dialog fields, click **Submit**. You have now created a backup. You can view this under the **Configuration** tab. The backup file will appear in the **Configuration** pane on the right.
- Note** When deployed and configured to do so, the Cisco Network Assurance Engine (NAE) creates export policies in the Cisco APIC for collecting data at timed intervals. You can identify a Cisco NAE export policy by its name, which is based on the assurance control configuration. If you delete a Cisco NAE export policy in the Cisco APIC, the Cisco NAE export policy will reappear in the Cisco APIC. We recommend that you do not delete the Cisco NAE export policies.
- Step 5** In the **Navigation** pane, choose **Export Policies > Configuration > policy_name**.
- Step 6** In the **Work** pane, choose the **Operational > Job Status** tabs.

On this screen, you can view a table with information about the export jobs. If you did not trigger an export job, then the table is empty. The **State** column indicates the status of an export job. The possible values are:

- **success**: The job succeeded.
- **failed**: The job failed.
- **success-with-warnings**: The job succeeded, but there were some issues.

The **Details** column indicates whether the integrity validation succeeded or failed.

If you created a backup, the Cisco APIC creates a file that is shown in the **Operational** view of the backup file that was created. If you want to then import that data, you must create an import policy.

Configuring an Import Policy Using the GUI

This procedure explains how to configure an Import policy using the APIC GUI. Follow these steps to import your backed up data:

Procedure

- Step 1** On the menu bar, choose **ADMIN > Import/Export**.
- Step 2** In the navigation pane, right-click **Import Policies** and click **Create Configuration Import Policy**. The **Create Configuration Import Policy** dialog appears.
- Step 3** Enter the appropriate values in the **Create Configuration Import Policy** dialog fields.
- Note** For an explanation of a field, click the 'i' icon to display the help file. For more detailed information on import types and modes including (**Replace**, **Merge**, **Best Effort**, and **Atomic**), refer to the *Cisco Application Centric Infrastructure Fundamentals Guide* .
- Step 4** When finished entering values in the **Create Configuration Import Policy** dialog fields, click **Submit**.
- Note** If you perform a clean reload of the fabric and import a previously-saved configuration, the time zone will change to UTC by default. Reset the time zone to your local time zone after the configuration import for the APIC cluster in these situations.

Encrypting Configuration Files Using the GUI

AES-256 encryption is a global configuration option. When enabled, all secure properties conform to the AES configuration setting. A portion of the ACI fabric configuration can be exported using configuration export with a specific targetDn. However, it is not possible to use REST API to export just a portion of the ACI fabric such as a tenant configuration with secure properties and AES encryption. The secure properties do not get included during REST API requests.

This section explains how to enable AES-256 encryption.

Procedure

-
- Step 1** On the menu bar, choose **ADMIN > AAA**.
- Step 2** In the navigation pane, click **AES Encryption Passphrase and Keys for Config Export (and Import)**. The **Global AES Encryption Settings for all Configurations Import and Export** window appears in the right pane.
- Step 3** Create a passphrase, which can be between 16 and 32 characters long. There are no restrictions on the type of characters used.
- Step 4** Click **SUBMIT**.
- Note** Once you have created and posted the passphrase, the keys are then generated in the back-end and the passphrase is not recoverable. Therefore, your passphrase is not visible to anyone because the key is automatically generated then deleted. Your backup only works if you know the passphrase (no one else can open it).
- The **Key Configured** field now shows **yes**. You now see an encrypted hash (which is not the actual passphrase, but just a hash of it) in the **Encrypted Passphrase** field.
- Step 5** After setting and confirming your passphrase, check the check box next to **Enable Encryption** to turn the AES encryption feature on (checked).
- The **Global AES Encryption Settings** field in your export and import policies will now be enabled by default.

- Note**
- Be sure that the **Fail Import if secure fields cannot be decrypted** check box is checked (which is the default selection) in your import and export policies. We highly recommend that you do not uncheck this box when you import configurations. If you uncheck this box, the system attempts to import all the fields. However, any fields that it cannot encrypt are blank/missing. As a result, you could lock yourself out of the system because the admin passwords could go blank/missing (if you lock yourself out of the system, refer to *Cisco APIC Troubleshooting Guide*). Unchecking the box launches a warning message. If the box is checked, there are security checks that prevent lockouts and the configuration does not import.
 - When the **Enable Encryption** check box is unchecked (off), encryption is disabled and all exported configurations (exports) are missing the secure fields (such as passwords and certificates). When this box is checked (on), encryption is enabled and all exports show the secure fields.
 - After enabling encryption, you cannot configure a passphrase when creating a new import or export policy. The passphrase you previously set is now global across all configurations in this box and across all tenants. If you export a configuration from this tab (you have configured a passphrase and enabled encryption) you get a complete backup file. If encryption is not enabled, you get a backup file with the secure properties removed. These backup files are useful when exporting to TAC support engineers, for example, because all the secure fields are missing. This is true for any secure properties in the configuration. There is also a clear option that clears the encryption key.

Note the list of the configuration import behaviors and associated results in the following table:

Configuration Import Behavior Scenario	Result
Old configuration from previous release	Import of configurations from old releases is fully supported and successfully imports all secure fields stored in old configurations.
Configuration import when AES encryption is not configured	If the import is for a configuration without secure fields, it is successful with the behavior previously described. If the imported configuration has secure fields, it is rejected.
Configuration import when AES passphrases do not match	If the import is for a configuration without secure fields, it is successful with the behavior previously described. If the imported configuration has secure fields, it is rejected.
Configuration import when AES passphrases match	Import is successful
Configuration import when AES passphrases do not match for copy/pasted fields	This specific case occurs when you have copied and pasted secure fields from other configurations that were exported with a different passphrase. During the first pass parsing of the imported backup file, if any property fails to decrypt correctly, the import fails without importing any shards. Therefore, if a shard fails to decrypt all properties, all shards are rejected.

Backing up, Restoring, and Rolling Back Controller Configuration

This section describes the set of features for backing up (creating snapshots), restoring, and rolling back a controller configuration.

Backing Up, Restoring, and Rolling Back Configuration Files Workflow

This section describes the workflow of the features for backing up, restoring, and rolling back configuration files. All of the features described in this document follow the same workflow pattern. Once the corresponding policy is configured, **adminSt** must be set to **triggered** in order to trigger the job.

Once triggered, an object of type **configJob** (representing that run) is created under a container object of type **configJobCont**. (The naming property value is set to the policy DN.) The container's **lastJobName** field can be used to determine the last job that was triggered for that policy.



Note Up to five **configJob** objects are kept under a single job container at a time, with each new job triggered. The oldest job is removed to ensure this.

The **configJob** object contains the following information:

- Execution time
- Name of the file being processed/generated
- Status, as follows:
 - Pending
 - Running
 - Failed
 - Fail-no-data
 - Success
 - Success-with-warnings
- Details string (failure messages and warnings)
- Progress percentage = $100 * \text{lastStepIndex} / \text{totalStepCount}$
- Field **lastStepDescr** indicating what was being done last

About the fileRemotePath Object

The fileRemotePath object holds the following remote location-path parameters:

- Hostname or IP
- Port
- Protocol: FTP, SCP, and others
- Remote directory (not file path)
- Username
- Password



Note The password must be resubmitted every time changes are made.

Sample Configuration

The following is a sample configuration:

Under **fabricInst** (uni/fabric), enter:

```
<fileRemotePath name="path-name" host="host name or ip" protocol="scp"
remotePath="path/to/some/folder" userName="user-name" userpasswd="password" />
```

Configuration Export to Controller

The configuration export extracts user-configurable managed object (MO) trees from all thirty-two shards in the cluster, writes them into separate files, then compresses them into a tar gzip. The configuration export then uploads the tar gzip to a pre-configured remote location (configured using **configRsRemotePath** pointing to a **fileRemotePath** object) or stores it as a **snapshot** on the controller(s).



Note See the Snapshots section for more details.

The **configExportP** policy is configured as follows:

- **name**: Policy name.
- **format**: Format in which the data is stored inside the exported archive (xml or json).
- **targetDn**: The domain name (DN) of the specific object you want to export (empty means everything).
- **snapshot**: When set to `True`, the file is stored on the controller, no remote location configuration is needed.
- **includeSecureFields**: Set to true by default, indicates whether the encrypted fields (passwords, etc.) should be included in the export archive.



Note The **configSnapshot** object is created holding the information about this snapshot (see the Snapshots section).

Scheduling Exports

An export policy can be linked with a scheduler, which triggers the export automatically based on a pre-configured schedule. This is done via the **configRsExportScheduler** relation from the policy to a **trigSchedP** object (see the following Sample Configuration section).



Note A scheduler is optional. A policy can be triggered at any time by setting the **adminSt** to **triggered**.

Troubleshooting

If you get an error message indicating that the generated archive could not be uploaded to the remote location, refer to the Connectivity Issues section.

Sample Configuration Using the NX-OS Style CLI

The following is a sample configuration using the NX-OS Style CLI:

```

apicl(config)# snapshot
download Configuration snapshot download setup mode
export Configuration export setup mode
import Configuration import setup mode
rollback Configuration rollback setup mode
upload Configuration snapshot upload setup mode
apicl(config)# snapshot export policy-name
apicl(config-export)#
format Snapshot format: xml or json
no Negate a command or set its defaults
remote Set the remote path configuration will get exported to
schedule Schedule snapshot export
target Snapshot target

bash bash shell for unix commands
end Exit to the exec mode
exit Exit from current mode
fabric show fabric related information
show Show running system information
where show the current mode
apicl(config-export)# format xml
apicl(config-export)# no remote path [If no remote path is specified, the file is
exported locally to a folder in the controller]
apicl(config-export)# target [Assigns the target of the export, which can
be fabric, infra, a specific tenant, or none. If no target is specified, all configuration
information is exported.]
WORD infra, fabric or tenant-x
apicl(config-export)#
apicl# trigger snapshot export policy-name [Executes the snapshot export task]
apicl# ls /data2 [If no remote path is specified, the
configuration export file is saved locally to the controller under the folder data2]
ce_Dailybackup.tgz

```

Sample Configuration Using the GUI

The following is a sample configuration using the GUI:

1. On the menu bar, click the **Admin** tab.
2. Choose **IMPORT/EXPORT**.

3. Under **Export Policies**, choose **Configuration**.
4. Under Configuration, click the configuration that you would like to roll back to. For example, you can click **defaultOneTime**, which is the default.
5. Next to **Format**, choose a button for either JSON or XML format.
6. Next to **Start Now**, choose a button for either **No** or **Yes** to indicate whether you want to trigger now or trigger based on a schedule. The easiest method is to choose to trigger immediately.
7. For the **Target DN** field, enter the name of the tenant configuration you are exporting.
8. If you want to store the configuration on the controller itself, check the **Snapshot** option. If you want to configure a remote location, uncheck this option.
9. For the **Scheduler** field, you have the option to create a scheduler instructing when and how often to export the configuration.
10. For the **Encryption** field, you have the option to enable or disable the encryption of your configuration file.
11. When you have finished your configuration, click **Start Now**.
12. Click **Submit** to trigger your configuration export.

Sample Configuration Using REST API

The following is a sample configuration using the REST API:

```
<configExportP name="policy-name" format="xml" targetDn="/some/dn or empty which means
everything"
snapshot="false" adminSt="triggered">
<configRsRemotePath tnFileRemotePathName="some remote path name" />
<configRsExportScheduler tnTrigSchedPName="some scheduler name" />
</configExportP>
```



Note When providing a remote location, if you set the snapshot to `True`, the backup ignores the remote path and stores the file on the controller.

Configuration Import to Controller

Configuration import downloads, extracts, parses, analyzes and applies the specified, previously exported archive one shard at a time in the following order: infra, fabric, tn-common, then everything else. The fileRemotePath configuration is performed the same way as for export (via configRsRemotePath). Importing snapshots is also supported.

The **configImportP** policy is configured as follows:

- **name** - policy name
- **fileName** - name of the archive file (not the path file) to be imported
- **importMode**

- Best-effort mode: each MO is applied individually, and errors only cause the invalid MOs to be skipped.



Note If the object is not present on the controller, none of the children of the object get configured. Best-effort mode attempts to configure the children of the object.

- Atomic mode: configuration is applied by whole shards. A single error causes whole shard to be rolled back to its original state.

- **importType**

- replace - Current system configuration is replaced with the contents or the archive being imported (only atomic mode is supported)
- merge - Nothing is deleted, archive content is applied on top the existing system configuration.

- **snapshot** - when true, the file is taken from the controller and no remote location configuration is needed.

- **failOnDecryptErrors** - (true by default) the file fails to import if the archive was encrypted with a different key than the one that is currently set up in the system.

Troubleshooting

The following scenarios may need troubleshooting:

- If the generated archive could not be downloaded from the remote location, refer to the Connectivity Issues section.
- If the import succeeded with warnings, check the details.
- If a file could not be parsed, refer to the following scenarios:
 - If the file is not a valid XML or JSON file, check whether or not the files from the exported archive were manually modified.
 - If an object property has an unknown property or property value, it may be because:
 - The property was removed or an unknown property value was manually entered
 - The model type range was modified (non-backward compatible model change)
 - The naming property list was modified
- If an MO could not be configured, note the following:
 - Best-effort mode logs the error and skips the MO
 - Atomic mode logs the error and skips the shard

Sample Configuration Using the NX-OS Style CLI

The following is a sample configuration using the NX-OS Style CLI:

```

apicl# configure
apicl(config)# snapshot
  download Configuration snapshot download setup mode
  export Configuration export setup mode
  import Configuration import setup mode
  rollback Configuration rollback setup mode
  upload Configuration snapshot upload setup mode
apicl(config)# snapshot import
  WORD Import configuration name
default
rest-user
apicl(config)# snapshot import policy-name
apicl(config-import)#
  action Snapshot import action merge|replace
  file Snapshot file name
  mode Snapshot import mode atomic|best-effort
  no Negate a command or set its defaults
  remote Set the remote path configuration will get imported from

bash bash shell for unix commands
end Exit to the exec mode
exit Exit from current mode
fabric show fabric related information
show Show running system information
where show the current mode
apicl(config-import)# file < from "show snapshot files" >
apicl(config-import)# no remote path
apicl(config-import)#
apicl# trigger snapshot import policy-name [Executes the snapshot import task]

```

Sample Configuration Using the GUI

The following is a sample configuration using the GUI:

1. On the menu bar, click the **ADMIN** tab.
2. Select **IMPORT/EXPORT**.
3. Under **Import Policies**, select **Configuration**.
4. Under **Configuration**, select **Create Configuration Import Policy**. The **CREATE CONFIGURATION IMPORT POLICY** window appears.
5. In the **Name** field, the file name must match whatever was backed up and will have a very specific format. The file name is known to whoever did the backup.
6. The next two options relate to recovering configuration states (also known as "roll-back"). The options are **Input Type** and **Input Mode**. When you recover a configuration state, you want to roll back to a known state that was good before. The option for that is an **Atomic Replace**.
7. If you want to store the configuration on the controller itself, check the **Snapshot** option. If you want to configure a remote location, uncheck this option.
8. In the **Import Source** field, specify the same remote location that you already created.
9. For the **Encryption** field, you have the option to enable or disable the encryption of your configuration file.
10. Click **SUBMIT** to trigger your configuration import.

Sample Configuration Using the REST API

The following shows a sample configuration using the REST API:

```
<configImportP name="policy-name" fileName="someexportfile.tgz" importMode="atomic"
importType="replace" snapshot="false" adminSt="triggered">
<configRsRemotePath tnFileRemotePathName="some remote path name" />
</configImportP>
```

Snapshots

Snapshots are configuration backup archives, stored (and replicated) in a controller managed folder. To create one, an export can be performed with the **snapshot** property set to true. In this case, no remote path configuration is needed. An object of **configSnapshot** type is created to expose the snapshot to the user.

You can create recurring snapshots, which are saved to **Admin > Import/Export > Export Policies > Configuration > defaultAuto**.

configSnapshot objects provide the following:

- file name
- file size
- creation date
- root DN indicating what the snapshot is of (fabric, infra, specific tenant, and so on)
- ability to remove a snapshot (by setting the retire field to true)

To import a snapshot, first create an import policy. Navigate to **Admin > Import/Export** and click **Import Policies**. Right click and choose **Create Configuration Import Policy** to set the import policy attributes.

Snapshot Manager Policy

The **configSnapshotManagerP** policy allows you to create snapshots from remotely stored export archives. You can attach a remote path to the policy, provide the file name (same as with configImportP), set the mode to download, and trigger. The manager downloads the file, analyzes it to make sure the archive is valid, stores it on the controller, and creates the corresponding configSnapshot object.

You can also create a recurring snapshot.



Note When enabled, recurring snapshots are saved to **Admin > Import/Export > Export Policies > Configuration > defaultAuto**.

The snapshot manager also allows you to upload a snapshot archive to a remote location. In this case, the mode must be set to upload.

Troubleshooting

For troubleshooting, refer to the Connectivity Issues section.

Snapshot Upload from Controller to Remote Path Using the NX-OS CLI

```

apicl(config)# snapshot upload policy-name
apicl(config-upload)#
  file      Snapshot file name
  no        Negate a command or set its defaults
  remote    Set the remote path configuration will get uploaded to

bash       bash shell for unix commands
end        Exit to the exec mode
exit       Exit from current mode
fabric     show fabric related information
show       Show running system information
where      show the current mode
apicl(config-upload)# file <file name from "show snapshot files">
apicl(config-upload)# remote path remote-path-name
apicl# trigger snapshot upload policy-name      [Executes the snapshot upload task]

```

Snapshot Download from Controller to Remote Path Using the NX-OS CLI

```

apicl(config)# snapshot download policy-name
apicl(config-download)#
  file      Snapshot file name
  no        Negate a command or set its defaults
  remote    Set the remote path configuration will get downloaded from

bash       bash shell for unix commands
end        Exit to the exec mode
exit       Exit from current mode
fabric     show fabric related information
show       Show running system information
where      show the current mode
apicl(config-download)# file < file from remote path>
apicl(config-download)# remote path remote-path-name
apicl# trigger snapshot download policy-name    [Executes the snapshot download task]

```

Snapshot Upload and Download Using the GUI

To upload a snapshot file to a remote location:

1. Right-click on the snapshot file listed in the **Config Rollbacks** pane, and select the **Upload to Remote Location option**. The **Upload snapshot to remote location** box appears.
2. Click **SUBMIT**.

To download a snapshot file from a remote location:

1. Click the import icon on the upper right side of the screen. The **Import remotely stored export archive to snapshot** box appears.
2. Enter the file name in the **File Name** field.
3. Select a remote location from the Import Source pull-down, or check the box next to **Or create a new one** to create a new remote location.
4. Click **SUBMIT**.

Snapshot Upload and Download Using the REST API

```

<configSnapshotManagerP name="policy-name" fileName="someexportfile.tgz"
mode="upload|download" adminSt="triggered">

```

```
<configRsRemotePath tnFileRemotePathName="some remote path name" />
</configSnapshotManagerP>
```

Rollback

The **configRollbackP** policy enables you to undo the changes made between two snapshots, effectively rolling back any configuration changes that were made to the snapshot that was saved earlier. When the policy is triggered, objects are processed as follows:

- Deleted MOs are recreated
- Created MOs are deleted
- Modified MOs are reverted



Note

- The rollback feature only operates on snapshots.
 - Remote archives are not supported directly. However, you can turn a remotely saved export into a snapshot using the snapshot manager policy (configSnapshotMgrP). For more information, see the [Snapshot Manager Policy, on page 25](#)
 - The configRollbackP policy does not require a remote path configuration. If a remote path is provided, it will be ignored.
-

Rollback Workflow

The policy snapshotOneDN and snapshotTwoDn fields must be set with the first snapshot (S1) preceding snapshot two (S2). When triggered, the snapshots are extracted and analyzed to calculate and apply the differences between the snapshots.

The MOs are handled as follows:

- MOs are present in S1 but not present in S2 — These MOs were deleted before S2. The rollback will recreate these MOs.
- MOs are present in S2 but not present in S1 — These MOs were created after S1. The rollback will delete these MOs under the following circumstances:
 - These MOs were not modified after S2 was taken.
 - No MO descendants were created or modified after S2 was taken.
- MOs are present in both S1 and S2 but with different property values — If the property was modified to a different value after S2 was taken, the property is left as is. Otherwise, the rollback will revert these properties to S1.

The rollback feature also generates a diff file that contains the configuration generated as a result of these calculations. Applying this configuration is the last step of the rollback process. The content of this file can be retrieved through a special REST API called readiff:

```
apichost/mqapi2/snapshots.readiff.xml?jobdn=SNAPSHOT_JOB_DN.
```

Rollback, which is difficult to predict, also has a preview mode (set `preview` to `true`), which prevents rollback from making any actual changes. It simply calculates and generates the diff file, allowing you to preview what exactly is going to happen once the rollback is actually performed.

Diff Tool

Another special REST API is available, which provides diff functionality between two snapshots:
`apichost/mqapi2/snapshots.diff.xml?s1dn=SNAPSHOT_ONE_DN&s2dn=SNAPSHOT_TWO_DN`.

Sample Configuration Using the NX-OS Style CLI

This example shows how to configure and execute a rollback using the NX-OS Style CLI:

```
apicl# show snapshot files
File      : ce2_DailyAutoBackup-2015-11-21T01-00-17.tar.gz
Created   : 2015-11-21T01:00:21.167+00:00
Root      :
Size      : 22926

File      : ce2_DailyAutoBackup-2015-11-21T09-00-21.tar.gz
Created   : 2015-11-21T09:00:24.025+00:00
Root      :
Size      : 23588

apicl# configure
apicl(config)# snapshot rollback myRollbackPolicy
apicl(config-rollback)# first-file ce2_DailyAutoBackup-2015-11-21T01-00-17.tar.gz
apicl(config-rollback)# second-file ce2_DailyAutoBackup-2015-11-21T09-00-21.tar.gz
apicl(config-rollback)# preview
apicl(config-rollback)# end
apicl# trigger snapshot rollback myRollbackPolicy
```

Sample Configuration Using the GUI

This example shows how to configure and execute a rollback using the GUI:

1. On the menu bar, click the **Admin** tab.
2. Click **Config Rollbacks**, located under the Admin tab.
3. Select the first configuration file from the **Config Rollbacks** list (in the left-side pane).
4. Select the second configuration file in the **Configuration for selected snapshot** pane (in the right-side pane).
5. Click the **Compare with previous snapshot** drop-down menu (at the bottom of the right-side pane), then select the second configuration file from that list. A diff file is then generated so that you can compare the differences between the two snapshots.



Note After the file generates, there is an option to undo these changes.

Sample Configuration Using the REST API

This example shows how to configure and execute a rollback using the REST API:

```
<configRollbackP name="policy-name" snapshotOneDn="dn/of/snapshot/one"  
snapshotOneDn="dn/of/snapshot/two" preview="false" adminSt="triggered" />
```

Using the Cisco APIC Troubleshooting Tools

This chapter introduces the tools and methodology commonly used to troubleshoot problems you may experience. These tools can assist you with monitoring traffic, debugging, and detecting issues such as traffic drops, misrouting, blocked paths, and uplink failures. See the tools listed below for a summary overview of the tools described in this chapter:

- **ACL Contract Permit and Deny Logs**—Enables the logging of packets or flows that were allowed to be sent because of contract permit rules and the logging of packets or flows dropped because of taboo contract deny rules.
- **Atomic Counters**—Enables you to gather statistics about traffic between flows for detecting drops and misrouting in the fabric and for enabling quick debugging and isolation of application connectivity issues.
- **Digital Optical Monitoring**—Enables you to view digital optical monitoring (DOM) statistics about a physical interface.
- **Health Scores**—Enables you to isolate performance issues by drilling down through the network hierarchy to isolate faults to specific managed objects (MOs).
- **Port Tracking**—Enables you to monitor the status of links between leaf switches and spine switches for detecting uplink failure.
- **SNMP**—Simple Network Management Protocol (SNMP) enables you to remotely monitor individual hosts (APIC or another host) and find out the state of any particular node.
- **SPAN**—Switchport Analyzer (SPAN) enables you to perform detailed troubleshooting or to take a sample of traffic from a particular application host for proactive monitoring and analysis.
- **Statistics**—Provides real-time measures of observed objects. Viewing statistics enable you to perform trend analysis and troubleshooting.
- **Syslog**—Enables you to specify the minimum severity level of messages to be sent, the items to be included in the syslog messages, and the syslog destination. The format can also be displayed in NX-OS CLI format.
- **Traceroute**—Enables you to find the routes that packets actually take when traveling to their destination.
- **Troubleshooting Wizard**—Enables administrators to troubleshoot issues that occur during specific time frames, which can be designated by selecting two endpoints.
- **Configuration Sync Issues**—Enables you to see if any transactions in Cisco APIC have not yet synced.

This chapter contains the following sections:

Using Atomic Counters

About Atomic Counters

Atomic counters allow you to gather statistics about traffic between flows. Using atomic counters, you can detect drops and misrouting in the fabric, enabling quick debugging and isolation of application connectivity

issues. For example, an administrator can enable atomic counters on all leaf switches to trace packets from endpoint 1 to endpoint 2. If any leaf switches have nonzero counters, other than the source and destination leaf switches, an administrator can drill down to those leafs.

In conventional settings, it is nearly impossible to monitor the amount of traffic from a bare metal NIC to a specific IP address (an endpoint) or to any IP address. Atomic counters allow an administrator to count the number of packets that are received from a bare metal endpoint without any interference to its data path. In addition, atomic counters can monitor per-protocol traffic that is sent to and from an endpoint or an application group.

Leaf-to-leaf (TEP-to-TEP) atomic counters can provide the following:

- Counts of sent, received, dropped, and excess packets
 - Sent packets: The sent number reflects how many packets were sent from the source TEP (tunnel endpoint) to the destination TEP.
 - Received packets: The received number reflects how many packets the destination TEP received from the source TEP.
 - Dropped packets: The dropped number reflects how many packets were dropped during transmission. This number is the difference in the amount of packets sent and the amount of packets received.
 - Excess packets: The excess number reflects how many extra packets were received during transmission. This number is the amount of packets that were unexpectedly received due to a forwarding mismatch or a misrouting to the wrong place.
- Short-term data collection such as the last 30 seconds, and long-term data collection such as 5 minutes, 15 minutes, or more
- A breakdown of per-spine traffic (available when the number of TEPs, leaf or VPC, is less than 64)
- Ongoing monitoring



Note Leaf-to-leaf (TEP to TEP) atomic counters are cumulative and cannot be cleared. However, because 30-second atomic counters reset at 30-second intervals, they can be used to isolate intermittent or recurring problems. Atomic counters require an active fabric Network Time Protocol (NTP) policy.

Tenant atomic counters can provide the following:

- Application-specific counters for traffic across the fabric, including sent, received, dropped, and excess packets
- Modes include the following:
 - EPtoEP (endpoint to endpoint)
 - EPGtoEPG (endpoint group to endpoint group)



Note For EPGtoEPG, the options include ipv4 only, ipv6 only, and ipv4, ipv6. Any time there is an ipv6 option, you use twice the TCAM entries, which means the scale numbers may be less than expected for pure ipv4 policies.

- EPGtoEP (endpoint group to endpoint)
- EPtoAny (endpoint to any)
- AnytoEP (any to endpoint)
- EPGtoIP (endpoint group to IP, used only for external IP address)
- EPtoExternalIP (endpoint to external IP address)

Beginning with the 5.2(3) release, endpoint security groups (ESGs) can be used as an alternative for EPGs in these modes.

Atomic Counters Guidelines and Restrictions

- Use of atomic counters is not supported when the endpoints are in different tenants or in different contexts (VRFs) within the same tenant.
- In Cisco APIC release 3.1(2m) and later, if no statistics have been generated on a path in the lifetime of the fabric, no atomic counters are generated for the path. Also, the **Traffic Map** in the **Visualization** tab (**Operations** > **Visualization** in the Cisco APIC GUI) does not show all paths, only the active paths (paths that had traffic at some point in the fabric lifetime).
- In pure Layer 2 configurations where the IP address is not learned (the IP address is 0.0.0.0), endpoint-to-EPG and EPG-to-endpoint atomic counter policies are not supported. In these cases, endpoint-to-endpoint and EPG-to-EPG policies are supported. External policies are virtual routing and forwarding (VRF)-based, requiring learned IP addresses, and are supported.
- When the atomic counter source or destination is an endpoint, the endpoint must be dynamic and not static. Unlike a dynamic endpoint (fv:CEp), a static endpoint (fv:StCEp) does not have a child object (fv:RsCEpToPathEp) that is required by the atomic counter.
- In a transit topology, where leaf switches are not in full mesh with all spine switches, then leaf-to-leaf (TEP to TEP) counters do not work as expected.
- For leaf-to-leaf (TEP to TEP) atomic counters, once the number of tunnels increases the hardware limit, the system changes the mode from trail mode to path mode and the user is no longer presented with per-spine traffic.
- The atomic counter does not count spine proxy traffic.
- Packets dropped before entering the fabric or before being forwarded to a leaf port are ignored by atomic counters.
- Packets that are switched in the hypervisor (same Port Group and Host) are not counted.
- Atomic counters require an active fabric Network Time Protocol (NTP) policy.
- Atomic counters work for IPv6 sources and destinations, but configuring source and destination IP addresses across IPv4 and IPv6 addresses is not allowed.
- An atomic counter policy configured with fvCEp as the source or destination counts only the traffic that is from/to the MAC and IP addresses that are present in the fvCEp managed objects. If the fvCEp managed object has an empty IP address field, then all traffic to/from that MAC address would be counted regardless of the IP address. If the Cisco APIC has learned multiple IP addresses for an fvCEp, then traffic from only the one IP address in the fvCEp managed object itself is counted as previously stated. To configure

an atomic counter policy to or from a specific IP address, use the fvIp managed object as the source or destination.

- If there is an fvIp behind an fvCEp, you must add fvIP-based policies and not fvCEp-based policies.
- Endpoint-to-endpoint atomic counter statistics are not reported for Layer 2 bridged traffic with IPv6 headers when the endpoints belong to the same EPG.
- For atomic counters to work for traffic flowing from an EPG or ESG to an L3Out EPG, configure the L3Out EPG with 0/1 and 128/1 to match all prefixes instead of 0/0.
- If your Cisco APIC has the traffic map mode set to "trial" and the Cisco APIC generated the F1545 fault, the only way that you can clear this fault is by setting the traffic map mode to "path." To change the traffic map mode, go to **Operations > Visualization**, click **Settings**, choose **path** for Mode, then click **Submit**. This will give you tunnel stats per port in both ingress and egress.

The trial mode has a greater chance of reaching the maximum scale index of tunnel logical interfaces. This mode consumes more software and hardware resources. A logical interface is the ID that is associated with the tunnel in the hardware.

If you have a single tunnel between a tunnel endpoint (TEP) you specified the trail mode, it will consume more hardware resources as well. For example, if you have 6 fabric ports and a single tunnel, then hardware consumes a number of entries equal to the number of tunnels multiplied by the number of fabric ports.

For software, if the number of logical interfaces allocated is greater than 2048, you will fail to have an entry in the hardware. As a result, you cannot get the stats. In the case of the atomic counter, this issue may show as drops or excesses.

The path mode has only entries for the TEP. For a vPC, two entries will be installed. Therefore, you have a lower chance of reaching to the maximum limit.

Configuring Atomic Counters

Procedure

-
- Step 1** In the menu bar, click **Tenants**.
- Step 2** In the submenu bar, click the desired tenant.
- Step 3** In the **Navigation** pane, expand the tenant and expand **Policies** and then expand **Troubleshoot**.
- Step 4** Under **Troubleshoot**, expand **Atomic Counter Policy** and choose a traffic topology.
You can measure traffic between a combination of endpoints, endpoint groups, external interfaces, and IP addresses.
- Step 5** Right-click the desired topology and choose **Add topology Policy** to open an **Add Policy** dialog box.
- Step 6** In the **Add Policy** dialog box, perform the following actions:
- a) In the **Name** field, enter a name for the policy.
 - b) choose or enter the identifying information for the traffic source.
The required identifying information differs depending on the type of source (endpoint, endpoint group, external interface, or IP address).
 - c) choose or enter the identifying information for the traffic destination.

d) (Optional) (Optional) In the **Filters** table, click the + icon to specify filtering of the traffic to be counted. In the resulting **Create Atomic Counter Filter** dialog box, you can specify filtering by the IP protocol number (TCP=6, for example) and by source and destination IP port numbers.

e) Click **Submit** to save the atomic counter policy.

Step 7 In the **Navigation** pane, under the selected topology, choose the new atomic counter policy. The policy configuration is displayed in the **Work** pane.

Step 8 In the **Work** pane, click the **Operational** tab and click the **Traffic** subtab to view the atomic counter statistics.

Enabling Atomic Counters

To enable using atomic counters to detect drops and misrouting in the fabric and enable quick debugging and isolation of application connectivity issues, create one or more tenant atomic counter policies, which can be one of the following types:

- EP_to_EP—Endpoint to endpoint (**dbgacEpToEp**)
- EP_to_EPG—Endpoint to endpoint group (**dbgacEpToEpg**)
- EP_to_Ext—Endpoint to external IP address (**dbgacEpToExt**)
- EPG_to_EP—Endpoint group to endpoint(**dbgacEpgToEp**)
- EPG_to_EPG—Endpoint group to endpoing group (**dbgacEpgToEpg**)
- EPG_to_IP—Endpoint group to IP address (**dbgacEpgToIp**)
- Ext_to_EP—External IP address to endpoint (**dbgacExtToEp**)
- IP_to_EPG—IP address to endpoint group (**dbgacIpToEpg**)
- Any_to_EP—Any to endpoint (**dbgacAnyToEp**)
- EP_to_Any—Endpoint to any (**dbgacEpToAny**)

Procedure

Step 1 To create an EP_to_EP policy using the REST API, use XML such as the following example:

Example:

```
<dbgacEpToEp name="EP_to_EP_Policy" ownerTag="" ownerKey=""
dn="uni/tn-Tenant64/acEpToEp-EP_to_EP_Policy" descr="" adminSt="enabled">
<dbgacFilter name="EP_to_EP_Filter" ownerTag="" ownerKey="" descr=""
srcPort="https" prot="tcp" dstPort="https"/>
</dbgacEpToEp>
```

Step 2 To create an EP_to_EPG policy using the REST API, use XML such as the following example:

Example:

```
<dbgacEpToEpg name="EP_to_EPG_Pol" ownerTag="" ownerKey=""
dn="uni/tn-Tenant64/epToEpg-EP_to_EPG_Pol" descr="" adminSt="enabled">
<dbgacFilter name="EP_to_EPG_Filter" ownerTag="" ownerKey="" descr=""
srcPort="http" prot="tcp" dstPort="http"/>
```

```
<dbgacRsToAbsEpg tDn="uni/tn-Tenant64/ap-VRF64_app_prof/epg-EPG64"/>
</dbgacEpToEpg>
```

Troubleshooting Using Atomic Counters with the REST API

Procedure

- Step 1** To get a list of the endpoint-to-endpoint atomic counters deployed within the fabric and the associated details such as dropped packet statistics and packet counts, use the **dbgEpToEpTsIt** class in XML such as the following example:

Example:

```
https://apic-ip-address/api/node/class/dbgEpToEpRslt.xml
```

- Step 2** To get a list of external IP-to-endpoint atomic counters and the associated details, use the **dbgacExtToEp** class in XML such as the following example:

Example:

```
https://apic-ip-address/api/node/class/dbgExtToEpRslt.xml
```

Enabling and Viewing Digital Optical Monitoring Statistics

Real-time digital optical monitoring (DOM) data is collected from SFPs, SFP+, and XFPs periodically and compared with warning and alarm threshold table values. The DOM data collected are transceiver transmit bias current, transceiver transmit power, transceiver receive power, and transceiver power supply voltage.

Enabling Digital Optical Monitoring Using the GUI

Before you can view digital optical monitoring (DOM) statistics about a physical interface, enable DOM on the leaf or spine interface, using a switch policy, associated to a policy group.

To enable DOM using the GUI:

Procedure

- Step 1** On the menu bar, choose **Fabric > Fabric Policies**.
- Step 2** In the **Navigation** pane, expand **Policies > Monitoring > Fabric Node Controls**.
- Step 3** Expand **Fabric Node Controls** to see a list of existing policies.
- Step 4** In the **Work** pane, click the **ACTIONS** drop-down menu and select **Create Fabric Node Control**. The **Create Fabric Node Control** dialog box appears.
- Step 5** In the **Create Fabric Node Control** dialog box, perform the following actions:
- In the **Name** field, enter a name for the policy.
 - Optional. In the **Description** field, enter a description of the policy.

c) Put a check in the box next to **Enable DOM**.

- Step 6** Click **Submit** to create the policy.
Now you can associate this policy to a policy group and a profile, as described in the following steps.
- Step 7** In the **Navigation** pane, expand **Switch Policies > Policy Groups**.
- Step 8** In the **Work** pane, click the **ACTIONS** drop-down menu and select **Create Leaf Switch Policy Group** (for a spine, **Create Spine Switch Policy Group**).
The **Create Leaf Switch Policy Group** or **Create Spine Switch Policy Group** dialog box appears.
- Step 9** In the dialog box, perform the following actions:
- In the **Name** field, enter a name for the policy group.
 - From the **Node Control Policy** drop-down menu, choose either an existing policy (such as the one you just created) or a new one by selecting **Create Fabric Node Control**.
 - Click **Submit**.
- Step 10** Attach the policy group you created to a switch as follows:
- In the **Navigation** pane, expand **Switch Policies > Profiles**.
 - In the **Work** pane, click the **ACTIONS** drop-down menu and select **Create Leaf Switch Profile** or **Create Spine Switch Profile**, as appropriate.
 - In the dialog box, enter a name for the profile in the **Name** field.
 - Add the name of the switch you want associated with the profile under **Switch Associations**.
 - From the **Blocks** pull-down menu, check the boxes next to the applicable switches.
 - From the **Policy Group** pull-down menu, select the policy group you created earlier.
 - Click **UPDATE**, then click **Submit**.

Enabling Digital Optical Monitoring Using the REST API

Before you can view digital optical monitoring (DOM) statistics about a physical interface, enable DOM on the interface.

To enable DOM using the REST API:

Procedure

- Step 1** Create a fabric node control policy (fabricNodeControlPolicy) as in the following example:

```
<fabricNodeControl dn="uni/fabric/nodecontrol-testdom" name="testdom" control="1"
rn="nodecontrol-testdom" status="created" />
```

- Step 2** Associate a fabric node control policy to a policy group as follows:

```
<?xml version="1.0" encoding="UTF-8" ?>
<fabricLeNodePGrp dn="uni/fabric/funcprof/lenodepgrp-nodegrp2" name="nodegrp2"
rn="lenodepgrp-nodegrp2" status="created,modified" >

    <fabricRsMonInstFabricPol tnMonFabricPolName="default" status="created,modified" />
    <fabricRsNodeCtrl tnFabricNodeControlName="testdom" status="created,modified" />

</fabricLeNodePGrp>
```

- Step 3** Associate a policy group to a switch (in the following example, the switch is 103) as follows:

```

<?xml version="1.0" encoding="UTF-8" ?>
<fabricLeafP>
  <attributes>
    <dn>uni/fabric/leprof-leafSwitchProfile</dn>
    <name>leafSwitchProfile</name>
    <rn>leprof-leafSwitchProfile</rn>
    <status>created,modified</status>
  </attributes>
  <children>
    <fabricLeafS>
      <attributes>
        <dn>uni/fabric/leprof-leafSwitchProfile/leaves-test-typ-range</dn>
        <type>range</type>
        <name>test</name>
        <rn>leaves-test-typ-range</rn>
        <status>created,modified</status>
      </attributes>
      <children>
        <fabricNodeBlk>
          <attributes>

<dn>uni/fabric/leprof-leafSwitchProfile/leaves-test-typ-range/nodeblk-09533c1d228097da</dn>

          <from_>103</from_>
          <to_>103</to_>
          <name>09533c1d228097da</name>
          <rn>nodeblk-09533c1d228097da</rn>
          <status>created,modified</status>
        </attributes>
        </fabricNodeBlk>
      </children>
    </children>
    <fabricRsLeNodePGrp>
      <attributes>
        <tDn>uni/fabric/funcprof/lenodepgrp-nodegrp2</tDn>
        <status>created</status>
      </attributes>
    </fabricRsLeNodePGrp>
  </children>
</fabricLeafS>
</children>
</fabricLeafP>

```

Viewing Digital Optical Monitoring Statistics With the GUI

To view DOM statistics using the GUI:

Before you begin

You must have previously enabled digital optical monitoring (DOM) statistics for an interface, before you can view the DOM statistics for it.

Procedure

- Step 1** In the Menu bar, choose **Fabric** and **Inventory**.

- Step 2** In the Navigation pane, expand the Pod and Leaf node where the physical interface you are investigating is located.
- Step 3** Expand **Interfaces**.
- Step 4** Expand **Physical Interfaces**.
- Step 5** Expand the physical interface you are investigating.
- Step 6** Choose **DOM Stats**.
DOM statistics are displayed for the interface.

Troubleshooting Using Digital Optical Monitoring With the REST API

To view DOM statistics using an XML REST API query:

Before you begin

You must have previously enabled digital optical monitoring (DOM) on an interface, before you can view the DOM statistics for it.

Procedure

The following example shows how to view DOM statistics on a physical interface, eth1/25 on node-104, using a REST API query:

```
GET
https://apic-ip-address/api/node/mo/topology/pod-1/node-104/sys/phys-[eth1/25]/phys/domstats.xml?
query-target=children&target-subtree-class=ethpmDOMRxPwrStats&subscription=yes
```

The following response is returned:

```
response : {
  "totalCount": "1",
  "subscriptionId": "72057611234705430",
  "imdata": [
    {"ethpmDOMRxPwrStats": {
      "attributes": {
        "alert": "none",
        "childAction": "",
        "dn": "topology/pod-1/node-104/sys/phys[eth1/25]/phys/domstats/rxpower",
        "hiAlarm": "0.158490",
        "hiWarn": "0.079430",
        "loAlarm": "0.001050",
        "loWarn": "0.002630",
        "modTs": "never",
        "status": "",
        "value": "0.139170" }}}}
```

Viewing and Understanding Health Scores

The APIC uses a policy model to combine data into a health score. Health scores can be aggregated for a variety of areas such as for infrastructure, applications, or services. The health scores enable you to isolate

performance issues by drilling down through the network hierarchy to isolate faults to specific managed objects (MOs). You can view network health by viewing the health of an application (by tenant) or by the health of a leaf switch (by pod).

For more information about health scores, faults, and health score calculation see the *Cisco APIC Fundamentals Guide*.

Health Score Types

The APIC supports the following health score types:

- System—Summarizes the health of the entire network.
- Leaf—Summarizes the health of leaf switches in the network. Leaf health includes hardware health of the switch including fan tray, power supply, and CPU.
- Tenant—Summarizes the health of a tenant and the tenant's applications.

Filtering by Health Score

You can filter health scores using the following tools:

- Health Scroll Bar—You can use the health scroll bar to dictate which objects are visible; lowering the score allows you to see only objects with a degraded health score.
- Displaying Degraded Health Scores—To display only the degraded health scores, click the Gear icon and choose **Show only degraded health score**.

Viewing Tenant Health

To view application health, click **Tenants** > *tenant-name* in the menu bar, then click the tenant name in the **Navigation** pane. The GUI displays a summary of the tenant's health including applications and EPGs. To drill down on the tenant configuration, double-click the health score.

For a health summary, click the **Health** tab in the **Work** pane. This view of the network displays health scores and relationships between MOs in the network so that you can isolate and resolve performance issues. For example, a common sequence of managed objects in the tenant context is **Tenant** > **Application profile** > **Application EPG** > **EPP** > **Fabric location** > **EPG to Path Attachment** > **Network Path Endpoint** > **Aggregation Interface** > **Aggregated Interface** > **Aggregated Member Interface**.

Viewing Fabric Health

To view fabric health, click **Fabric** in the menu bar. In the **navigation** pane, choose a pod. The GUI displays a summary of the pod health including nodes. To drill down on part of the fabric configuration, double-click the health score.

For a health summary, click the **Health** tab in the **work** pane. This view of the network displays health scores and relationships between MOs in the network so that you can isolate and resolve performance issues. For example, a common sequence of managed objects in the fabric context is **Pod** > **Leaf** > **Chassis** > **Fan tray slot** > **Line module slot** > **Line module** > **Fabric Port** > **Layer 1 Physical Interface Configuration** > **Physical Interface Runtime State**.



Note Fabric issues, such as physical network problems, can impact tenant performance when MOs are directly related.

Viewing MO Health in Visore

To view the health of an MO in Visore, click the **H** icon.

Use the following MOs to display health information:

- health:Inst
- health:NodeInst
- observer:Node
- observer:Pod

For more information about Visore, see the *Cisco Application Centric Infrastructure Fundamentals* guide.

Debugging Health Scores Using Logs

You can use the following log files to debug health scores on the APIC:

- svc_ifc_eventmgr.log
- svc_ifc_observer.log

Check the following items when debugging health scores using logs:

- Verify the source of the syslog (fault or event).
- Check whether a syslog policy is configured on the APIC.
- Check whether the syslog policy type and severity is set correctly.
- You can specify a syslog destination of console, file, RemoteDest, or Prof. ForRemoteDest, ensure that the syslog server is running and reachable.

Viewing Faults

The steps below explain where to view fault information.

Procedure

- Step 1** Go to a faults window:
- System Faults—From the menu bar, click **System > Faults**.
 - Tenant Faults—From the menu bar:
 - a. Click **Tenants > tenant-name**.
 - b. From the **Navigation** pane, click the **Tenants** *tenant name*.

- c. From the **Work** pane, click the **Faults** tab.
- Fabric Faults—From the menu bar:
 - a. Click **Fabric > Inventory**.
 - b. From the **Navigation** pane, click on a **Pod**
 - c. From the **Work** pane, click the **Faults** tab.

A list of faults appears in a summary table.

Step 2 Double-click on a fault.

The fabric and system tables change to display faults that match the fault code of the fault you clicked on.

- a) From the fabric or system faults, double-click on a fault in the summary table to view more information.

The **Fault Properties** dialog appears displaying the following tabs:

- **General**—Displays the following:
 - **Properties**—Contains information found in the summary table
 - **Details**—Contains fault information found in the summary table, the number of occurrences, the change set, and the original, previous, and highest severity level for the chosen fault.
- **Troubleshooting**—Displays the following:
 - **Troubleshooting**—Contains troubleshooting information that includes an explanation of the fault and the recommended action.
 - **Audit log**—A tool that enables you to view the history of user-initiated events before the fault occurred. The history is displayed in a list by a specified number of minutes. You can adjust the number of minutes by clicking the drop-down arrow.
- **History**—Displays history information of the affected object

Enabling Port Tracking for Uplink Failure Detection

This section explains how to enable port tracking using the GUI, NX-OS CLI, and the REST API.

Port Tracking Policy for Fabric Port Failure Detection

Fabric port failure detection can be enabled in the port tracking system settings. The port tracking policy monitors the status of fabric ports between leaf switches and spine switches, and ports between tier-1 leaf switches and tier-2 leaf switches. When an enabled port tracking policy is triggered, the leaf switches take down all access interfaces on the switch that have EPGs deployed on them.

If you enabled the **Include APIC ports when port tracking is triggered** option, port tracking disables Cisco Application Policy Infrastructure Controller (APIC) ports when the leaf switch loses connectivity to all fabric ports (that is, there are 0 fabric ports). Enable this feature only if the Cisco APICs are dual- or multihomed

to the fabric. Bringing down the Cisco APIC ports helps in switching over to the secondary port in the case of a dual-homed Cisco APIC.



Note Port tracking is located under **System > System Settings > Port Tracking**.

The port tracking policy specifies the number of fabric port connections that trigger the policy, and a delay timer for bringing the leaf switch access ports back up after the number of specified fabric ports is exceeded.

The following example illustrates how a port tracking policy behaves:

- The port tracking policy specifies that the threshold of active fabric port connections each leaf switch that triggers the policy is 2.
- The port tracking policy triggers when the number of active fabric port connections from the leaf switch to the spine switches drops to 2.
- Each leaf switch monitors its fabric port connections and triggers the port tracking policy according to the threshold specified in the policy.
- When the fabric port connections come back up, the leaf switch waits for the delay timer to expire before bringing its access ports back up. This gives the fabric time to reconverge before allowing traffic to resume on leaf switch access ports. Large fabrics may need the delay timer to be set for a longer time.



Note Use caution when configuring this policy. If the port tracking setting for the number of active spine ports that triggers port tracking is too high, all leaf switch access ports will be brought down.

Configuring Port Tracking Using the GUI

This procedure explains how to use the Port Tracking feature using the GUI.

Procedure

-
- Step 1** From the **System** menu, select **System Settings**.
 - Step 2** In the navigation pane, select **Port Tracking**.
 - Step 3** Turn on the Port Tracking feature by selecting **on** next to **Port tracking state**.
 - Step 4** Turn off the Port Tracking feature by selecting **off** next to Port tracking state under Properties.
 - Step 5** (Optional) Reset the **Delay restore timer** from the default (120 seconds).
 - Step 6** Enter the maximum number of active spine links (any configuration value from 0 - 12) that are up before port tracking is triggered.
 - Step 7** Click **Submit** to push your desired Port Tracking configuration to all switches on the fabric.
-

Port Tracking Using the NX-OS CLI

This procedure explains how to use the Port Tracking feature using the NX-OS CLI.

Procedure

Step 1 Turn on the Port Tracking feature as follows:

Example:

```
apicl# show porttrack
Configuration
Admin State           : on
Bringup Delay(s)     : 120
Bringdown # Fabric Links up : 0
```

Step 2 Turn off the Port Tracking feature as follows:

Example:

```
apicl# show porttrack
Configuration
Admin State           : off
Bringup Delay(s)     : 120
Bringdown # Fabric Links up : 0
```

Port Tracking Using the REST API

Before you begin

This procedure explains how to use the Port Tracking feature using the REST API.

Procedure

Step 1 Turn on the Port Tracking feature using the REST API as follows (**admin state: on**):

```
<polUni>
<infraInfra dn="uni/infra">
<infraPortTrackPol name="default" delay="5" minlinks="4" adminSt="on">

</infraPortTrackPol>
</infraInfra>
</polUni>
```

Step 2 Turn off the Port Tracking feature using the REST API as follows (**admin state: off**):

```
<polUni>
<infraInfra dn="uni/infra">
<infraPortTrackPol name="default" delay="5" minlinks="4" adminSt="off">

</infraPortTrackPol>
</infraInfra>
</polUni>
```

Using SNMP

About SNMP

The Cisco Application Centric Infrastructure (ACI) provides extensive SNMPv1, v2, and v3 support, including Management Information Bases (MIBs) and notifications (traps). The SNMP standard allows any third-party applications that support the different MIBs to manage and monitor the Cisco ACI fabric.

SNMPv3 provides extended security. Each SNMPv3 device can be selectively enabled or disabled for SNMP service. In addition, each device can be configured with a method of handling SNMPv1 and v2 requests.

Beginning in the 4.2(6) release, SNMPv3 supports the Secure Hash Algorithm-2 (SHA-2) authentication type.

For more information about using SNMP, see the *Cisco ACI MIB Quick Reference*.

SNMP Access Support in Cisco ACI



Note For the complete list of MIBs supported in Cisco Application Centric Infrastructure (ACI), see <http://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/mib/list/mib-support.html>.

SNMP support in Cisco ACI is as follows:

- SNMP read queries (Get, Next, Bulk, Walk) are supported by leaf and spine switches and by the Cisco Application Policy Infrastructure Controller (APIC).
- SNMP write commands (Set) are not supported by leaf and spine switches or by the Cisco APIC.
- SNMP traps (v1, v2c, and v3) are supported by leaf and spine switches and by the Cisco APIC.



Note Cisco ACI supports a maximum of 10 trap receivers.

- SNMPv3 is supported by leaf and spine switches and by the Cisco APIC.
- SNMP using a Cisco APIC IPv6 address is not supported.

Table 1: SNMP Support Changes by Cisco APIC Release

Release	Description
1.2(2)	IPv6 support is added for SNMP trap destinations.
1.2(1)	SNMP support for the Cisco APIC controller is added. Previous releases support SNMP only for leaf and spine switches.

SNMP Trap Aggregation

The SNMP Trap Aggregation feature allows SNMP traps from the fabric nodes to be aggregated by Cisco Application Policy Infrastructure Controllers (APICs) and allows the forwarding of SNMP traps received from the fabric nodes to the external destination by the APICs.

Use this feature if you expect traps to come from APIC instead of from individual fabric nodes. With this feature enabled, APIC acts as an SNMP proxy.

We highly recommend that you configure all APICs in the cluster as SNMP trap aggregators to handle possible failures. You can configure multiple trap destinations in the SNMP policy. To configure trap aggregation and forwarding, follow these steps:

1. Configure each APIC controller to receive traps from the switches. Follow the procedure in [Configuring an SNMP Trap Destination Using the GUI, on page 46](#) using the following settings:
 - In the **Host Name/IP** field, specify the IPv4 or IPv6 address of the APIC.
 - From the **Management EPG** list, select the out-of-band or inband management EPG.

Repeat this procedure to configure each APIC in the cluster as a trap destination.

2. Configure the APIC to forward aggregated traps to an external server. Follow the procedure in [Configuring the SNMP Policy Using the GUI, on page 44](#) using the following settings:
 - In the **Trap Forward Servers** table, add the **IP Address** of the external server.

With trap aggregation and forwarding, the source IP address of the forwarded trap is the address of the aggregator (in this case, the APIC) and not the actual source node. To determine the actual source, you must search in the OID. In the following example, the address 10.202.0.1 is the APIC IP address, and the address 10.202.0.201 is the IP address of the original source leaf switch.

```
08:53:10.372378 IP
(tos 0x0, ttl 60, id 59067, offset 0, flags [DF], proto UDP (17), length 300)
10.202.0.1.45419 > 192.168.254.200.162: [udp sum ok]
{ SNMPv2c C="SNMP-ACI" { V2Trap(252) R=609795065
.1.3.6.1.2.1.1.3.0=25847714 .1.3.6.1.6.3.1.1.4.1.0=.1.3.6.1.4.1.9.9.276.0.1
.1.3.6.1.2.1.2.2.1.1.436207616=436207616 .1.3.6.1.2.1.2.2.1.7.436207616=2
.1.3.6.1.2.1.2.2.1.8.436207616=2 .1.3.6.1.2.1.31.1.1.1.1.436207616="eth1/1"
.1.3.6.1.2.1.2.2.1.3.436207616=6 .1.3.6.1.2.1.2.2.1.2.436207616="eth1/1"
.1.3.6.1.2.1.31.1.1.1.18.436207616=""
.1.3.6.1.4.1.9.10.22.1.4.1.1.6="10.202.0.201" } }
```

The SNMP Trap Aggregation feature was introduced in the Cisco APIC release 3.1(1) with support for SNMPV2 trap aggregation and forwarding. Beginning in the Cisco APIC releases 4.2(6) and 5.1(1), SNMPv3 trap aggregation and forwarding is supported.



Note If an APIC is decommissioned, the user is expected to clean reboot the decommissioned APIC. Since SNMP Trap Aggregation functionality is active on decommissioned APICs, the user could receive duplicate traps on the trap destination if the decommissioned APIC is not clean rebooted.

Configuring SNMP

Configuring the SNMP Policy Using the GUI

This procedure configures and enables the SNMP policy on Cisco ACI-mode switches.

Before you begin

To allow SNMP communications, you must configure the following:

- Configure an out-of-band contract allowing SNMP traffic. SNMP traffic typically uses UDP port 161 for SNMP requests.
- Configure the Cisco APIC out-of-band IP addresses in the "mgmt" tenant. If you configured the out-of-band addresses during the Cisco APIC setup, do not the configure the addresses in the "mgmt" tenant again.

Procedure

Step 1 In the menu bar, click **Fabric**.

Step 2 In the submenu bar, click **Fabric Policies**.

Step 3 In the **Navigation** pane, expand **Pod Policies**.

Step 4 Under **Pod Policies**, expand **Policies**.

Step 5 Right-click **SNMP** and choose **Create SNMP Policy**.

As an alternative to creating a new SNMP policy, you can edit the **default** policy fields in the same manner as described in the following steps.

Step 6 In the SNMP policy dialog box, perform the following actions:

- a) In the **Name** field, enter an SNMP policy name.
- b) In the **Admin State** field, select **Enabled**.
- c) (Optional) In the **SNMP v3 Users** table, click the + icon, enter a **Name**, enter the user's authentication data, and click **Update**.

This step is needed only if SNMPv3 access is required.

- d) In the **Community Policies** table, click the + icon, enter a **Name**, and click **Update**.

The community policy name can be a maximum of 32 characters in length. The name can contain only letters, numbers and the special characters of underscore (_), hyphen (-), or period (.). The name cannot contain the @ symbol.

- e) In the **Trap Forward Servers** table, click the + icon, enter the **IP Address** of the external server and click **Update**.

Step 7 Required: To configure allowed SNMP management stations, perform the following actions in the SNMP policy dialog box:

- a) In the **Client Group Policies** table, click the + icon to open the **Create SNMP Client Group Profile** dialog box.
- b) In the **Name** field, enter an SNMP client group profile name.
- c) From the **Associated Management EPG** drop-down list, choose the management EPG.
- d) In the **Client Entries** table, click the + icon.
- e) Enter a client's name in the **Name** field, enter the client's IP address in the **Address** field, and click **Update**.

Note When an SNMP management station connects with APIC using SNMPv3, APIC does not enforce the client IP address specified in the SNMP client group profile. For SNMPv3, the management station must exist in the **Client Entries** list, but the IP address need not match, as the SNMPv3 credentials alone are sufficient for access.

- Step 8** Click **OK**.
- Step 9** Click **Submit**.
- Step 10** Under **Pod Policies**, expand **Policy Groups** and choose a policy group or right-click **Policy Groups** and choose **Create POD Policy Group**.
- You can create a new pod policy group or you can use an existing group. The pod policy group can contain other pod policies in addition to the SNMP policy.
- Step 11** In the pod policy group dialog box, perform the following actions:
- In the **Name** field, enter a pod policy group name.
 - From the **SNMP Policy** drop-down list, choose the SNMP policy that you configured and click **Submit**.
- Step 12** Under **Pod Policies**, expand **Profiles** and click **default**.
- Step 13** In the **Work pane**, from the **Fabric Policy Group** drop-down list, choose the pod policy group that you created.
- Step 14** Click **Submit**.
- Step 15** Click **OK**.

Configuring an SNMP Trap Destination Using the GUI

This procedure configures the host information for an SNMP manager that will receive SNMP trap notifications.



Note ACI supports a maximum of 10 trap receivers. If you configure more than 10, some will not receive notifications.

Procedure

- Step 1** In the menu bar, click **Admin**.
- Step 2** In the submenu bar, click **External Data Collectors**.
- Step 3** In the **Navigation** pane, expand **Monitoring Destinations**.
- Step 4** Right-click **SNMP** and choose **Create SNMP Monitoring Destination Group**.
- Step 5** In the **Create SNMP Monitoring Destination Group** dialog box, perform the following actions:
- In the **Name** field, enter an SNMP destination name and click **Next**.
 - In the **Create Destinations** table, click the + icon to open the **Create SNMP Trap Destination** dialog box.
 - In the **Host Name/IP** field, enter an IPv4 or IPv6 address or a fully qualified domain name for the destination host.
 - Choose the **Port** number and **SNMP Version** for the destination.
 - For SNMP v1 or v2c destinations, enter one of the configured community names as the **Security Name** and choose **noauth** as **v3 Security Level**.

An SNMP v1 or v2c security name can be a maximum of 32 characters in length. The name can contain only letters, numbers and the special characters of underscore (_), hyphen (-), or period (.). For SNMP v1, the name cannot contain the @ symbol.

For SNMP v2c, in the 4.2(6) release and earlier, the name cannot contain the @ symbol. In the 4.2(7) release and later, the name can contain the @ symbol.

- f) For SNMP v3 destinations, enter one of the configured SNMP v3 user names as **Security Name** and choose the desired **v3 Security Level**.

An SNMP v3 security name can be a maximum of 32 characters in length. The name must begin with an uppercase or lowercase letter, and can contain only letters, numbers, and the special characters of underscore (_), hyphen (-), or period (.). In the 4.2(6) release and earlier, the name cannot contain the @ symbol. In the 4.2(7) release and later, the name can contain the @ symbol.

- g) From the **Management EPG** drop-down list, choose the management EPG.
h) Click **OK**.
i) Click **Finish**.

Configuring an SNMP Trap Source Using the GUI

This procedure selects and enables a source object within the fabric to generate SNMP trap notifications.

Procedure

-
- Step 1** In the menu bar, click **Fabric**.
- Step 2** In the submenu bar, click **Fabric Policies**.
- Step 3** In the **Navigation** pane, expand **Monitoring Policies**.
You can create an SNMP source in the **Common Policy**, the **default** policy, or you can create a new monitoring policy.
- Step 4** Expand the desired monitoring policy and choose **Callhome/SNMP/Syslog**.
If you chose the **Common Policy**, right-click **Common Policy**, choose **Create SNMP Source**, and follow the instructions below for that dialog box.
- Step 5** In the **Work** pane, from the **Monitoring Object** drop-down list, choose **ALL**.
- Step 6** From the **Source Type** drop-down list, choose **SNMP**.
- Step 7** In the table, click the + icon to open the **Create SNMP Source** dialog box.
- Step 8** In the **Create SNMP Source** dialog box, perform the following actions:
- In the **Name** field, enter an SNMP policy name.
 - From the **Dest Group** drop-down list, choose an existing destination for sending notifications or choose **Create SNMP Monitoring Destination Group** to create a new destination.
The steps for creating an SNMP destination group are described in a separate procedure.
 - Click **Submit**.

Monitoring the System Using SNMP

You can remotely monitor individual hosts (APIC or another host) and find out the state of any particular node.

You can check the system's CPU and memory usage using SNMP to find out if the CPU is spiking or not. The SNMP, a network management system, uses an SNMP client and accesses information over the APIC and retrieves information back from it.

You can remotely access the system to figure out if the information is in the context of the network management system and you can learn whether or not it is taking too much CPU or memory, or if there are any system or performance issues. Once you learn the source of the issue, you can check the system health and verify whether or not it is using too much memory or CPU.

Refer to the *Cisco ACI MIB Quick Reference Manual* for additional information.

Using SPAN

About SPAN

You can use the Switched Port Analyzer (SPAN) utility to perform detailed troubleshooting or to take a sample of traffic from a particular application host for proactive monitoring and analysis.

SPAN copies traffic from one or more ports, VLANs, or endpoint groups (EPGs) and sends the copied traffic to one or more destinations for analysis by a network analyzer. The process is nondisruptive to any connected devices and is facilitated in the hardware, which prevents any unnecessary CPU load.

You can configure SPAN sessions to monitor traffic received by the source (ingress traffic), traffic transmitted from the source (egress traffic), or both. By default, SPAN monitors all traffic, but you can configure filters to monitor only selected traffic.

You can configure SPAN on a tenant or on a switch. When configured on a switch, you can configure SPAN as a fabric policy or an access policy.

APIC supports the encapsulated remote extension of SPAN (ERSPAN).

Beginning with Release 4.1(1i), the following features are now supported:

- Support for local SPAN with static port-channels as the destination, as long as the sources and the port-channel are local on the same switch.



Note If you are running APIC release 4.1(1i) or later and you configure a static port-channel as the destination, but then downgrade to a release prior to 4.1(1i), then the SPAN session will go into the administrator disabled state because this feature was not available prior to release 4.1.(1i). There is no other functionality impact.

- You no longer have to include the IP prefix of the Layer 3 interface when configuring source SPAN with Layer 3 interface filtering.
- Support for configuring filter groups, which is a grouping of one or more filter entries. Use the filter group to specify the matching criteria that will be used to determine if a received packet should be analyzed using SPAN.
- The SPAN-on-drop feature, which captures packets that are dropped due to forwarding at the ingress in the ASIC and sends them to a pre-configured SPAN destination. There are 3 types of SPAN-on-drop configuration: access drop using access ports as a SPAN source, fabric drop using fabric ports as a SPAN source, and global drop using all ports on a node as a SPAN source. SPAN-on-drop is configured using

regular SPAN (through the CLI, GUI, and REST API) and using troubleshooting SPAN (CLI and REST API, only). For more information about configuring this feature, see *Configuring SPAN Using the GUI*, *Configuring SPAN Using the NX-OS Style CLI*, and *Configuring SPAN Using the REST API*.

Multinode SPAN

The APIC traffic monitoring policies can span policies at the appropriate places to keep track of all the members of each application group and where they are connected. If any member moves, the APIC automatically pushes the policy to the new leaf. For example, when an endpoint VMotions to a new leaf, the span configuration automatically adjusts.

The ACI fabric supports the following two extensions of encapsulated remote SPAN (ERSPAN) formats:

- Access or tenant SPAN—done for leaf switch front panel ports with or without using VLAN as a filter. The Broadcom Trident 2 ASIC in the leaf switches supports a slightly different version of the ERSPAN Type 1 format. It differs from the ERSPAN Type 1 format defined in the document referenced above in that the GRE header is only 4 bytes and there is no sequence field. The GRE header is always encoded with the following – 0x000088be. Even though 0x88be indicates ERSPAN Type 2, the remaining 2 bytes of the fields identify this as an ERSPAN Type 1 packet with a GRE header of 4 bytes.
- Fabric SPAN—done in leaf switches by the Northstar ASIC or by the Alpine ASIC in the spine switches. While these ASICs support ERSPAN Type 2 and 3 formats, the ACI fabric currently only supports ERSPAN Type 2 for fabric SPAN, as documented in the base-line document referenced above.

Refer to the IETF Internet Draft at the following URL for descriptions of ERSPAN headers: <https://tools.ietf.org/html/draft-foschiano-erspan-00>.

SPAN Guidelines and Restrictions



Note Many guidelines and restrictions depend on whether the switch is a generation 1 or generation 2 switch. The generation of the switch is defined as follows:

- Generation 1 switches are identified by the lack of a suffix, such as "EX", "FX", or "FX2," at the end of the switch name (for example, N9K-9312TX).
 - Generation 2 switches are identified with a suffix, such as "EX", "FX", or "FX2," at the end of the switch name.
-
- The type of SPAN supported varies:
 - For generation 1 switches, tenant and access SPAN use the encapsulated remote extension of SPAN (ERSPAN) type I (Version 1 option in the Cisco Application Policy Infrastructure Controller (APIC) GUI).
 - For generation 2 switches, tenant and access SPAN use the encapsulated remote extension of SPAN (ERSPAN) type II (Version 2 option in the Cisco APIC GUI).
 - Fabric SPAN uses ERSPAN type II.
 - Beginning with the 5.2(3) release, ERSPAN supports IPv6 destinations.
 - The 6.0(3) release supports the Cisco N9K-C9808 switch, which has the following SPAN limitations:

- Egress (transit (Tx)) SPAN is not supported.
- SPAN on drop is not supported.
- You cannot have the same SPAN sources in multiple sessions.
- SPAN supports an MTU of up to 343 bytes.
- A uSeg EPG or ESG cannot be used as a SPAN source EPG because the SPAN source filter is based on the VLAN ID. Thus, even if an endpoint is classified to a uSeg EPG or an ESG, traffic from the endpoint is mirrored if its VLAN is the VLAN of the SPAN source EPG.
- When configuring ERSPAN session, if the SPAN source contains a destination and interfaces from a spine switch within a GOLF VRF instance, an L3Out prefix is sent to the GOLF router with the wrong BGP next-hop, breaking connectivity from GOLF to that L3Out.
- You cannot specify an l3extLifP Layer 3 subinterface as a SPAN source. You must use the entire port for monitoring traffic from external sources.
- In local SPAN for FEX interfaces, the FEX interfaces can only be used as SPAN sources, not SPAN destinations.
 - On generation 1 switches, Tx SPAN does not work for any Layer 3 switched traffic.
 - On generation 2 switches, Tx SPAN does not work whether traffic is Layer 2 or Layer 3 switched.

There are no limitations for Rx SPAN.

For SPAN of FEX fabric port-channel (NIF), the member interfaces are supported as SPAN source interfaces on generation 1 leaf switches.



Note While you can also configure FEX fabric port-channel (NIF) member interfaces as SPAN source interfaces on generation 2 switches, this is not supported for releases prior to Cisco APIC release 4.1.

For information regarding ERSPAN headers, refer to the IETF Internet Draft at this URL:
<https://tools.ietf.org/html/draft-foschiano-erspan-00>.

- ERSPAN destination IP addresses must be learned in the fabric as an endpoint.
- SPAN supports IPv6 traffic.
- The individual port member of a port channel or a vPC cannot be configured as the source. Use the port channel, vPC, or vPC component as the source in the SPAN session.
- A fault is not raised on the ERSPAN source group when the destination EPG is deleted or unavailable.
- SPAN filters are supported on generation 2 leaf switches only.

An access SPAN source supports only one of the following filters at a given time:

- EPG
- Routed outside (L3Out)

- When deploying the access SPAN source with an L3Out filter, ensure that the L3Out is also deployed on the matching interface:
 - If an L3Out is deployed on a port, a SPAN source must be deployed on the same port.
 - If an L3Out is deployed on a PC, a SPAN source must be deployed on the same PC.
 - If an L3Out is deployed on a vPC, a SPAN source must be deployed on the same vPC.
- An L3Out routed interface and routed sub-interface can be deployed on a port or a PC, but an L3Out SVI can be deployed on a port, PC, or vPC. A SPAN source with an L3Out filter must be deployed accordingly.
- An L3Out filter is not supported in fabric SPAN or tenant SPAN sessions.
- The correct L3Out must be selected in the L3 configuration tab of the EPG bridge domain; otherwise, packet flow for basic L3Out will not work.
- An encapsulation value is mandatory for a routed sub-interface and SVI, but is not applicable for a routed interface. The L3Out sub-interface or SVI encapsulation value must be different from the EPG encapsulation value.

When an EPG filter is enabled within a SPAN session, ARP packets, which are sent out of the interface in the transit, or tx, direction, will not be spanned.
- SPAN filters are not supported in the following:
 - Fabric ports
 - Fabric and tenant SPAN sessions
 - Spine switches
- L4 port range filter entries will not be added if you attempt to add more L4 port ranges than are officially supported.
- A SPAN session will not come up if you attempt to associate more than the supported filter entries at the SPAN source group level or at the individual SPAN source level.
- Deleted filter entries will remain in TCAM if you add or delete more filters entries than are officially supported.
- See the *Verified Scalability Guide for Cisco ACI* document for SPAN-related limits, such as the maximum number of active SPAN sessions and SPAN filter limitations.
- For the SPAN-on-drop feature, the following guidelines and restrictions apply:
 - The SPAN-on-drop feature is supported on generation 2 leaf switches.
 - The SPAN-on-drop feature only captures packets with forwarding drops in the LUX block, which captures forwarding drop packets at the ingress. The SPAN-on-drop feature cannot capture the BMX (buffer) and RWX (egress) drops.
 - When using the troubleshooting CLI to create a SPAN session with SPAN-on-drop enabled and Cisco APIC as the destination, the session is disabled when 100 MB of data is captured.
 - On a modular chassis, the SPAN-on-drop feature will only work for the packets dropped on the line cards. Packets that are dropped on the fabric card will not be spanned.

- SPAN-on-drop ACLs with other SPAN ACLs are not merged. If a SPAN-on-drop session is configured on an interface along with ACL-based SPAN, then any packets dropped on that interface will only be sent to the SPAN-on-drop session.
- You cannot configure SPAN on drop and SPAN ACL on the same session.
- When an access or fabric port-drop session and a global-drop session are configured, the access or fabric port-drop session takes the priority over the global-drop session.
- The number of filter entries supported in TCAM = $(M * S1 * 1 + N * S2 * 2) + (S3 * 2)$. This is applicable to rx SPAN or tx SPAN, separately. Currently, the maximum filter entries supported in tx or rx SPAN is 480 in each direction when following this formula (and assuming there are no other sources that are configured without filter-group association [means $S3 = 0$] and with 16 port-ranges included). When the number of filter entries exceed the maximum number allowed, a fault will be raised. Note that you can specify Layer 4 port ranges in the filter entry. However, sixteen Layer 4 ports are programmed into the hardware as a single filter entry.

**Note**

- M=The number of IPv4 filters
- S1=The number of sources with IPv4 filters
- N=The number of IPv6 filters
- S2=The number of sources with IPv6 filters
- S3=The number of sources with no filter group association

- With MAC pinning configured in the LACP policy for a PC or vPC, the PC member ports will be placed in the LACP individual port mode and the PC is operationally non-existent. Hence, a SPAN source configuration with such a PC will fail, resulting in the generation of the "No operational src/dst" fault. With the MAC pinning mode configured, SPAN can be configured only on individual ports.
- A packet that is received on a Cisco Application Centric Infrastructure (ACI) leaf switch will be spanned only once, even if span sessions are configured on both the ingress and egress interfaces.
- When you use a routed outside SPAN source filter, you see only unicast in the Tx direction. In the Rx direction, you can see unicast, broadcast, and multicast.
- An L3Out filter is not supported for transmit multicast SPAN. An L3Out is represented as a combination of sclass/dclass in the ingress ACL filters and can therefore match unicast traffic only. Transmit multicast traffic can be spanned only on ports and port-channels.
- You can use a port channel interface as a SPAN destination only on -EX and later switches.
- You cannot configure multiple SPAN sessions with the same source interface when a SPAN filter (5-tuple filter) is applied.

The local SPAN destination port of a leaf switch does not expect incoming traffic. You can ensure that the switch drops incoming SPAN destination port traffic by configuring a Layer 2 interface policy and setting the **VLAN Scope** property to **Port Local scope** instead of **Global scope**. Apply this policy to the SPAN destination ports. You can configure an Layer 2 interface policy by going to the following location in the GUI: **Fabric > Access Policies > Policies > Interface > L2 Interface**.

When you configure SPAN for a given packet, SPAN is supported for the packet only once. If traffic is selected by SPAN in Rx for the first SSN, the traffic will not be selected by SPAN again in Tx for a second SSN. Thus, when the SPAN session ingress and egress ports sit on a single switch, the SPAN session capture will be one-way only. The SPAN session cannot display two-way traffic.

- A SPAN ACL filter configured in the filter group does not filter the broadcast, unknown-unicast and multicast (BUM) traffic that egresses the access interface. A SPAN ACL in the egress direction works only for unicast IPv4 or IPv6 traffic.

When configuring a SPAN destination as a local port, EPGs cannot be deployed to that interface.

In a leaf switch, a SPAN source with a VRF filter will match all regular bridge domains and all Layer 3 SVIs under the VRF instance.

In a spine switch, a SPAN source with a VRF matches only the configured VRF VNID traffic and a bridge domain filter will match only the bridge domain VNID traffic.

- When you create your own SPAN extended filter entries, you cannot use `__UI_AUTO_CONFIG_DEFAULT_EXTENDED_MO` as an object name to identify your extended filter entry managed object.
- Use SPAN destination interfaces that have the same speed. Traffic monitored by a SPAN session may face traffic loss due to SPAN buffer drops even though the destination port is not oversubscribed, but one of the other SPAN destination ports is oversubscribed. The SPAN traffic rate is limited to the slowest SPAN destination interface speed when the destination interfaces have with different speeds and when one of them is oversubscribed.
- Use a SPAN destination interface speed that is higher than any configured source interfaces and choose a speed that contains enough leeway for micro-bursts. Cloud Scale ASICs do not provide a micro-burst monitoring option for the SPAN class.

Configuring SPAN Using the GUI

Configuring a Tenant SPAN Session Using the Cisco APIC GUI

SPAN can be configured on a switch or on a tenant. This section guides you through the Cisco APIC GUI to configure a SPAN policy on a tenant to forward replicated source packets to a remote traffic analyzer. The configuration procedure requires entering values in the fields of one or more GUI dialog boxes. To understand a field and determine a valid value, view the help file by clicking the help icon (?) at the top-right corner of the dialog box.

Procedure

-
- Step 1** In the menu bar, click **Tenants**.
 - Step 2** In the submenu bar, click the tenant that contains the source endpoint.
 - Step 3** In the **Navigation** pane, expand the tenant, expand **Policies > Troubleshooting > SPAN**.
Two nodes appear under **SPAN**: **SPAN Destination Groups** and **SPAN Source Groups**.
 - Step 4** From the **Navigation** pane, right-click **SPAN Source Groups** and choose **Create SPAN Source Group**. The **Create SPAN Source Group** dialog appears.
 - Step 5** Enter the appropriate values in the required fields of the **Create SPAN Source Group** dialog box.

- Step 6** Expand the **Create Sources** table to open the **Create SPAN Source** dialog box.
- Step 7** Enter the appropriate values in the **Create SPAN Source** dialog box fields.
- Step 8** When finished creating the SPAN source, click **OK**.
You return to the **Create SPAN Source Group** dialog box.
- Step 9** When finished entering values in the **Create SPAN Source Group** dialog box fields, click **Submit**.

What to do next

Using a traffic analyzer at the SPAN destination, you can observe the data packets from the SPAN source EPG to verify the packet format, addresses, protocols, and other information.

Configuring a SPAN Filter Group Using the APIC GUI

Procedure

-
- Step 1** In the menu bar, click on **Fabric** and in the submenu bar click on **Access Policies**.
- Step 2** In the **Navigation** pane, expand **Policies > Troubleshooting**, and expand **SPAN**.
- Step 3** Under **SPAN**, right-click **SPAN Filter Groups** and choose **Create SPAN Filter Group**. The **Create Filter Group** dialog appears.
- Step 4** Enter a name for the SPAN filter group. In the **Filter Entries** table, click + and enter values for the following fields:
- **Source IP Prefix:** Enter a source IP address in the form of *IP-address/mask*. Both IPv4 and IPv6 addresses are supported. Use a value of **0.0.0.0** to denote an IPv4 address **any** entry in this field or use a value of **::** to denote an IPv6 address **any** entry in this field.
 - **First Source Port:** Enter the first source Layer 4 port. This field, together with the **Last Source Port** field, specifies a port range for filtering source ports. Use a value of **0** to denote an **any** entry in this field.
 - **Last Source Port:** Enter the last source Layer 4 port. This field, together with the **First Source Port** field, specifies a port range for filtering source ports. Use a value of **0** to denote an **any** entry in this field.
 - **Destination IP Prefix:** Enter a destination IP address in the form of *IP-address/mask*. Both IPv4 and IPv6 addresses are supported. Use a value of **0.0.0.0** to denote an IPv4 address **any** entry in this field or use a value of **::** to denote an IPv6 address **any** entry in this field.
 - **First Destination Port:** Enter the first destination Layer 4 port. This field, together with the **Last Destination Port** field, specifies a port range for filtering destination ports. Use a value of **0** to denote an **any** entry in this field.
 - **Last Destination Port:** Enter the last destination Layer 4 port. This field, together with the **First Destination Port** field, specifies a port range for filtering destination ports. Use a value of **0** to denote an **any** entry in this field.
 - **IP Protocol:** Enter the IP protocol. Use a value of **0** to denote an **any** entry in this field.
 - In the **Extended Filter Entries** table, click + and enter values for the following fields:
 - **Name:** Enter a name for the extended filter entry.

- **DSCP From:** Enter the DSCP value. This field, together with the **DSCP To** field, specifies the range for filtering DSCP values.
- **DSCP To:** Enter the DSCP value. This field, together with the **DSCP From** field, specifies the range for filtering DSCP values.
- **Dot1P From:** Enter the Dot1P value. This field, together with the **Dot1P To** field, specifies the range for filtering Dot1P values.
- **Dot1P To:** Enter the Dot1P value. This field, together with the **Dot1P From** field, specifies the range for filtering Dot1P values.

You can either specify the values for the source port and the destination port range or the DSCP and Dot1P range. If you specify both the source port and the destination port range and the DSCP and Dot1P range, faults are displayed.

DSCP or Dot1P is not supported for the egress direction. If you choose **Both** as the direction, then either DSCP or Dot1P is supported for ingress direction only and not for egress direction.

- **TCP Flags:** Choose a protocol from the **TCP Flags** drop-down list.

You can configure **TCP Flags** only if you chose **Unspecified** or **TCP** as the **IP Protocol** in the drop-down list for the filter group.

- **Packet Type:** Choose the packet type. You can either choose **Routed/Switched**, **Routed**, or **Switched Only**.

- Step 5** Click **Update**, then click **Submit** when you have entered the appropriate values into each of the fields in this form.

Configuring an Access SPAN Policy Using the Cisco APIC GUI

This procedure guides you through the Cisco APIC GUI to configure an access SPAN policy. The configuration procedure requires entering values in the fields of one or more GUI dialog boxes.

Procedure

-
- Step 1** In the menu bar, click on **Fabric > Access Policies**.
- Step 2** In the **Navigation** pane, expand **Policies > Troubleshooting > SPAN**.
Three nodes appear under **SPAN**: **SPAN Source Groups**, **SPAN Filter Groups**, and **SPAN Destination Groups**.
- Step 3** Right-click **SPAN Source Groups** and choose **Create SPAN Source Group**.
The **Create SPAN Source Group** dialog appears.
- Step 4** Enter the appropriate values in the **Create SPAN Source Group** dialog box fields.
- Step 5** Expand the **Create Sources** table to open the **Create SPAN Source** dialog box and enter the appropriate values in the required fields.
- Step 6** In the **Create SPAN Source** dialog box, expand **Add Source Access Paths** to specify the source path.
The **Associate Source to Path** dialog box appears.

- Step 7** Enter the appropriate values in the **Associate Source to Path** dialog box fields.
- Step 8** When finished associating the source to a path, click **OK**.
You return to the **Create SPAN Source** dialog box.
- Step 9** When finished configuring the SPAN source, click **OK**.
You return to the **Create SPAN Source Group** dialog box.
- Step 10** When finished configuring the SPAN source group, click **Submit**.
-

What to do next

Using a traffic analyzer at the SPAN destination, you can observe the data packets from the SPAN source to verify the packet format, addresses, protocols, and other information.

Configuring a Fabric SPAN Policy Using the Cisco APIC GUI

This section guides you through the Cisco APIC GUI to create a fabric SPAN policy. The configuration procedure requires entering values in the fields of one or more GUI dialog boxes.

Procedure

- Step 1** In the menu bar, click on **Fabric > Fabric Policies**.
- Step 2** In the **Navigation** pane, expand **Policies > Troubleshooting > SPAN**.
Three nodes appear under **SPAN**: **SPAN Source Groups**, **SPAN Filter Groups**, and **SPAN Destination Groups**.
- Step 3** Right-click **SPAN Source Groups** and choose **Create SPAN Source Group**.
The **Create SPAN Source Group** dialog appears.
- Step 4** Enter the appropriate values in the **Create SPAN Source Group** dialog box fields.
- Step 5** Expand the **Create Sources** table to open the **Create SPAN Source** dialog box.
- Step 6** Enter the appropriate values in the **Create SPAN Source** dialog box fields.
- Step 7** When finished, click **OK**.
You return to the **Create SPAN Source Group** dialog box.
- Step 8** When finished entering values in the **Create SPAN Source Group** dialog box fields, click **Submit**.
-

What to do next

Using a traffic analyzer at the SPAN destination, you can observe the data packets from the SPAN source to verify the packet format, addresses, protocols, and other information.

Configuring a Layer 3 EPG SPAN Session for External Access Using the APIC GUI

This procedure shows how to configure a Layer 3 EPG SPAN policy for External Access using the Cisco APIC GUI. The configuration procedure requires entering values in the fields of one or more GUI dialog boxes.

Procedure

-
- Step 1** In the menu bar, click on **Fabric > Access Policies**.
- Step 2** In the **Navigation** pane, expand **Policies > Troubleshooting > SPAN**.
- Three nodes appear under **SPAN**: **SPAN Source Groups**, **SPAN Filter Groups**, and **SPAN Destination Groups**.
- Step 3** Right-click **SPAN Source Groups** and choose **Create SPAN Source Group**. The **Create SPAN Source Group** dialog appears.
- Step 4** Enter the appropriate values in the **Create SPAN Source Group** dialog box fields.
- Step 5** In the **Filter Group** field, select or create a filter group.
- See [Configuring a SPAN Filter Group Using the APIC GUI, on page 54](#) for more information.
- Step 6** Expand the **Create Sources** table to open the **Create SPAN Source** dialog box and perform the following actions:
- Enter a **Name** for the source policy.
 - Choose a **Direction** option for the traffic flow.
 - (Optional) Click to place a check mark in the **Span Drop Packets** check box. When checked, the SPAN-on-drop feature is enabled.
 - For external access, click **Routed Outside** in the **Type** field.
- Note** If **Routed Outside** is chosen for external access, then the **Name**, **Address**, and **Encap** fields appear to configure the **L3 Outside**.
- Expand **Add Source Access Paths** to specify the source path.
- The **Associate Source to Path** dialog box appears.
- Enter the appropriate values in the **Associate Source to Path** dialog box fields.
 - When finished associating the source to a path, click **OK**.
- You return to the **Create SPAN Source** dialog box.
- When finished configuring the SPAN source, click **OK**.
- You return to the **Create SPAN Source Group** dialog box.
- Step 7** When finished configuring the SPAN source group, click **Submit**.
-

What to do next

Using a traffic analyzer at the SPAN destination, you can observe the data packets from the SPAN source to verify the packet format, addresses, protocols, and other information.

Configuring a Destination Group for an Access SPAN Policy Using the Cisco APIC GUI

This section guides you through the Cisco APIC GUI to create a destination group for an access SPAN policy. The configuration procedure requires entering values in the fields of one or more GUI dialog boxes.

Creating a SPAN destination group and source enables you to use a traffic analyzer at the SPAN destination to observe the data packets from the SPAN source and verify the packet format, addresses, protocols, and other information.

Procedure

-
- Step 1** In the menu bar, click on **Fabric > Access Policies**.
 - Step 2** In the **Navigation** pane, expand **Policies > Troubleshooting > SPAN**.
Three nodes appear under **SPAN**: **SPAN Source Groups**, **SPAN Filter Groups**, and **SPAN Destination Groups**.
 - Step 3** Right-click **SPAN Destination Groups** and choose **Create SPAN Destination Group**.
The **Create SPAN Destination Group** dialog appears.
 - Step 4** Enter the appropriate values in the **Create SPAN Destination Group** dialog box fields.
 - Step 5** When finished, click **Submit**.
The destination group is created.
-

Configuring a Destination Group for a Fabric SPAN Policy Using the Cisco APIC GUI

This section guides you through the Cisco APIC GUI to create a destination group for a fabric SPAN policy. The configuration procedure requires entering values in the fields of one or more GUI dialog boxes.

Creating a SPAN destination group and source enables you to use a traffic analyzer at the SPAN destination to observe the data packets from the SPAN source and verify the packet format, addresses, protocols, and other information.

Procedure

-
- Step 1** In the menu bar, click on **Fabric > Fabric Policies**.
 - Step 2** In the **Navigation** pane, expand **Policies > Troubleshooting > SPAN**.
Three nodes appear under **SPAN**: **SPAN Source Groups**, **SPAN Filter Groups**, and **SPAN Destination Groups**.
 - Step 3** Right-click **SPAN Destination Groups** and choose **Create SPAN Destination Group**.
The **Create SPAN Destination Group** dialog appears.
 - Step 4** Enter the appropriate values in the **Create SPAN Destination Group** dialog box fields.
 - Step 5** When finished, click **Submit**.
The destination group is created.
-

What to do next

If not already created, configure a source for the fabric SPAN policy.

Configuring SPAN Using the NX-OS-Style CLI**Configuring Local SPAN in Access Mode Using the NX-OS-Style CLI**

This is the traditional SPAN configuration local to an Access leaf node. Traffic originating from one or more access ports or port-channels can be monitored and sent to a destination port local to the same leaf node.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: apicl# configure terminal	Enters global configuration mode.
Step 2	[no] monitor access session <i>session-name</i> Example: apicl(config)# monitor access session mySession	Creates an access monitoring session configuration.
Step 3	[no] description <i>text</i> Example: apicl(config-monitor-access)# description "This is my SPAN session"	Adds a description for this access monitoring session. If the text includes spaces, it must be enclosed in single quotes.
Step 4	[no] destination interface ethernet <i>slot/port leaf node-id</i> Example: apicl(config-monitor-access)# destination interface ethernet 1/2 leaf 101	Specifies the destination interface. The destination interface cannot be a FEX port.
Step 5	[no] source interface ethernet {[fex]/slot/port port-range} leaf <i>node-id</i> Example: apicl(config-monitor-access)# source interface ethernet 1/2 leaf 101	Specifies the source interface port or port range.
Step 6	drop enable Example: apicl(config-monitor-access-source)# drop enable	Enables the SPAN on drop feature, which captures all packets that are dropped in the ASIC and sends them to a pre-configured SPAN destination.

	Command or Action	Purpose
Step 7	<p>[no] direction {rx tx both}</p> <p>Example:</p> <pre>apicl(config-monitor-access-source)# direction tx</pre>	Specifies direction of traffic to be monitored. The direction can be configured independently for each source port range.
Step 8	<p>[no] filter tenant <i>tenant-name</i> application <i>application-name</i> epg <i>epg-name</i></p> <p>Example:</p> <pre>apicl(config-monitor-access-source)# filter tenant t1 application appl epg epg1</pre>	Filters traffic to be monitored. The filter can be configured independently for each source port range.
Step 9	<p>exit</p> <p>Example:</p> <pre>apicl(config-monitor-access-source)# exit</pre>	Returns to access monitor session configuration mode.
Step 10	<p>[no] destination interface port-channel <i>port-channel-name-list</i> leaf <i>node-id</i></p> <p>Example:</p> <pre>apicl(config-monitor-access)# destination interface port-channel pc1 leaf 101</pre>	Specifies the destination interface. The destination interface cannot be a FEX port. Note Beginning with Release 4.1(1), support is now available for having a static port-channel as the destination interface, as shown in the example command.
Step 11	<p>[no] source interface port-channel <i>port-channel-name-list</i> leaf <i>node-id</i> [fex <i>fex-id</i>]</p> <p>Example:</p> <pre>apicl(config-monitor-access)# source interface port-channel pc5 leaf 101</pre>	Specifies the source interface port channel. (Enters the traffic direction and filter configuration, not shown here.)
Step 12	<p>[no] filter tenant <i>tenant-name</i> l3out <i>L3Out-name</i> vlan <i>interface-VLAN</i></p> <p>Example:</p> <pre>apicl(config-monitor-access-source)# filter tenant t1 l3out l3out1 vlan 2820</pre>	Filters traffic to be monitored. The filter can be configured independently for each source port range. Note Beginning with Release 4.1(1), you no longer have to specify the IP prefix when configuring L3Out interface filtering, as shown in the example.
Step 13	<p>[no] shutdown</p> <p>Example:</p> <pre>apicl(config-monitor-access)# no shut</pre>	Disables (or enables) the monitoring session.

Examples

This example shows how to configure a local access monitoring session.

```

apic1# configure terminal
apic1(config)# monitor access session mySession
apic1(config-monitor-access)# description "This is my SPAN session"
apic1(config-monitor-access)# destination interface ethernet 1/2 leaf 101
apic1(config-monitor-access)# source interface ethernet 1/1 leaf 101
apic1(config-monitor-access)# drop enable
apic1(config-monitor-access-source)# direction tx
apic1(config-monitor-access-source)# filter tenant t1 application appl epg epg1
apic1(config-monitor-access-source)# exit
apic1(config-monitor-access)# no shut
apic1(config-monitor-access)# show run
# Command: show running-config monitor access session mySession
# Time: Fri Nov 6 23:55:35 2015
monitor access session mySession
  description "This is my SPAN session"
  destination interface eth 1/2 leaf 101
  source interface eth 1/1 leaf 101
  direction tx
  filter tenant t1 application appl epg epg
  exit
exit

```

Configuring a SPAN Filter Group Using the NX-OS-Style CLI

These procedures describe how to configure a SPAN filter group and filter entries.

Procedure

-
- | | |
|---------------|--|
| Step 1 | <p>configure</p> <p>Enters global configuration mode.</p> <p>Example:</p> <pre>apic1# configure</pre> |
| Step 2 | <p>[no] monitor access filter-group <i>filtergroup-name</i></p> <p>Creates an access monitoring filter group configuration.</p> <p>Example:</p> <pre>apic1(config)# monitor access filter-group filtergroup1</pre> |
| Step 3 | <p>[no] filter srcaddress <i>source-address</i> dstaddress <i>destination-address</i> srcport-from <i>source-from-port</i> srcport-to <i>source-to-port</i> dstport-from <i>destination-from-port</i> dstport-to <i>destination-to-port</i> ipproto <i>IP-protocol</i></p> <p>Configures the filter entries for the filter group, where:</p> <ul style="list-style-type: none"> • <i>source-address</i> is a source IP address in the form of <i>IP-address/mask</i>. Both IPv4 and IPv6 addresses are supported. Use a value of 0.0.0.0 to denote an IPv4 address any entry in this field or use a value of :: to denote an IPv6 address any entry in this field. |

- *destination-address* is a destination IP address in the form of *IP-address/mask*. Both IPv4 and IPv6 addresses are supported. Use a value of **0.0.0.0** to denote an IPv4 address **any** entry in this field or use a value of **::** to denote an IPv6 address **any** entry in this field.
- *source-from-port* is the first source Layer 4 port. This field, together with the **srcport-to** field, specifies a port range for filtering source ports. Use a value of **0** to denote an **any** entry in this field.
- *source-to-port* is the last source Layer 4 port. This field, together with the **srcport-from** field, specifies a port range for filtering source ports. Use a value of **0** to denote an **any** entry in this field.
- *destination-from-port* is the first destination Layer 4 port. This field, together with the **dstport-to** field, specifies a port range for filtering destination ports. Use a value of **0** to denote an **any** entry in this field.
- *destination-to-port* is the last destination Layer 4 port. This field, together with the **dstport-from** field, specifies a port range for filtering destination ports. Use a value of **0** to denote an **any** entry in this field.
- *IP-protocol* is the IP protocol. Use a value of **0** to denote an **any** entry in this field.

Example:

```
apicl(config-monitor-fltgrp)# filter srcaddress 1.1.1.0/24 dstaddress 0.0.0.0 srcport-from
 0 srcport-to 0 dstport-from 0 dstport-to 0 ipproto 20
```

Step 4 **exit**

Returns to access monitor filter group configuration mode.

Example:

```
apicl(config-monitor-fltgrp)# exit
```

Step 5 **exit**

Exits global configuration mode.

Example:

```
apicl(config)# exit
```

Examples

This example shows how to configure a SPAN filter group and filter entries.

```
apicl# configure
apicl(config)# monitor access filter-group filtergroup1
apicl(config-monitor-fltgrp)# filter srcaddress 1.1.1.0/24 dstaddress 0.0.0.0 srcport-from
 0 srcport-to 0 dstport-from 0 dstport-to 0 ipproto 20
apicl(config-monitor-fltgrp)# exit
apicl(config)# exit
```

Configuring a SPAN Filter With Extended Filters Using the NX-OS-Style CLI

The following example shows you how to configure a SPAN filter and the extended filters using the CLI.

Procedure

To configure a SPAN filter and the extended filters using the CLI:

Example:

```
apic1(config-monitor-access-filtergrp-filter-extended-filters)# show run
# Command: show running-config monitor access filter-group filtergroup1 filter dstaddr
192.168.10.1 srcaddr 192.168.10.100 extended-filters ext1
# Time: Wed May 11 11:25:23 2022
  monitor access filter-group filtergroup1
    filter srcaddr 192.168.10.100 dstaddr 192.168.10.1
      extended-filters ext1
        dscp from CS0 to 4
        dot1p from 1 to 5
        forwarding-type switched
        tcp-flag ack off
        tcp-flag fin off
        tcp-flag rst on
      exit
    exit
  exit
apic1#
```

Associating a SPAN Filter Group Using the NX-OS-Style CLI

These procedures describe how to associate a filter group to a SPAN source group.

Procedure

Step 1 **configure**

Enters global configuration mode.

Example:

```
apic1# configure
```

Step 2 **[no] monitor access session *session-name***

Creates an access monitoring session configuration.

Example:

```
apic1(config)# monitor access session session1
```

Step 3 **filter-group *filtergroup-name***

Associates a filter group.

Example:

```
apic1(config-monitor-access)# filter-group filtergroup1
```

Step 4 **no filter-group**

Disassociates a filter group, if necessary.

Example:

```
apicl(config-monitor-access)# no filter-group
```

Step 5 [no] **source interface ethernet** {[*flex*]/*slot/port* | *port-range*} **leaf** *node-id*

Specifies the source interface port or port range.

Example:

```
apicl(config-monitor-access)# source interface ethernet 1/9 leaf 101
```

Step 6 **filter-group** *filtergroup-name*

Associates a filter group to a SPAN source.

Example:

```
apicl(config-monitor-access-source)# filter-group filtergroup2
```

Step 7 **exit**

Returns to access monitor filter group configuration mode.

Example:

```
apicl(config-monitor-access-source)# exit
```

Step 8 **no filter-group**

Disassociates the filter group from a SPAN source, if necessary.

Example:

```
apicl(config-monitor-access-source)# no filter-group
```

Step 9 **exit**

Returns to access monitor filter group configuration mode.

Example:

```
apicl(config-monitor-access)# exit
```

Step 10 **exit**

Exits global configuration mode.

Example:

```
apicl(config)# exit
```

Examples

This example shows how to associate a filter group.

```
apicl# configure
apicl(config)# monitor access session session1
apicl(config-monitor-access)# filter-group filtergroup1
apicl(config-monitor-access)# source interface ethernet 1/9 leaf 101
apicl(config-monitor-access-source)# filter-group filtergroup2
apicl(config-monitor-access-source)# exit
apicl(config-monitor-access-source)# no filter-group
```



```
apic1(config-monitor-access)# exit
apic1(config)# exit
```

Configuring ERSPAN in Access Mode Using the NX-OS-Style CLI

In the ACI fabric, an access mode ERSPAN configuration can be used for monitoring traffic originating from access ports, port-channels, and vPCs in one or more leaf nodes.

For an ERSPAN session, the destination is always an endpoint group (EPG) which can be deployed anywhere in the fabric. The monitored traffic is forwarded to the destination wherever the EPG is moved.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: apic1# configure terminal	Enters global configuration mode.
Step 2	[no] monitor access session <i>session-name</i> Example: apic1(config)# monitor access session mySession	Creates an access monitoring session configuration.
Step 3	[no] description <i>text</i> Example: apic1(config-monitor-access)# description "This is my access ERSPAN session"	Adds a description for this monitoring session. If the text includes spaces, it must be enclosed in single quotes.
Step 4	[no] destination tenant <i>tenant-name</i> application <i>application-name</i> epg <i>epg-name</i> destination-ip <i>dest-ip-address</i> source-ip-prefix <i>src-ip-address</i> Example: apic1(config-monitor-access)# destination tenant t1 application appl1 epg epg1 destination-ip 192.0.20.123 source-ip-prefix 10.0.20.1	Specifies the destination interface as a tenant and enters destination configuration mode.
Step 5	[no] erspan-id <i>flow-id</i> Example: apic1(config-monitor-access-dest)# erspan-id 100	Configures the ERSPAN ID for the ERSPAN session. The ERSPAN range is from 1 to 1023.
Step 6	[no] ip dscp <i>dscp-code</i> Example: apic1(config-monitor-access-dest)# ip dscp 42	Configures the differentiated services code point (DSCP) value of the packets in the ERSPAN traffic. The range is from 0 to 64.

	Command or Action	Purpose
Step 7	[no] ip ttl <i>ttl-value</i> Example: apicl(config-monitor-access-dest) # ip ttl 16	Configures the IP time-to-live (TTL) value for the ERSPAN traffic. The range is from 1 to 255.
Step 8	[no] mtu <i>mtu-value</i> Example: apicl(config-monitor-access-dest) # mtu 9216	Configures the maximum transmit unit (MTU) size for the ERSPAN session. The range is 64 to 9216 bytes.
Step 9	exit Example: apicl(config-monitor-access-dest) #	Returns to monitor access configuration mode.
Step 10	[no] source interface ethernet {[<i>fex</i>]/ <i>slot/port</i> <i>port-range</i> } leaf <i>node-id</i> Example: apicl(config-monitor-access) # source interface eth 1/2 leaf 101	Specifies the source interface port or port range.
Step 11	[no] source interface port-channel <i>port-channel-name-list</i> leaf <i>node-id</i> [fex <i>fex-id</i>] Example: apicl(config-monitor-access) # source interface port-channel pc1 leaf 101	Specifies the source interface port-channel.
Step 12	[no] source interface vpc <i>vpc-name-list</i> leaf <i>node-id1</i> <i>node-id2</i> [fex <i>fex-id1</i> <i>fex-id2</i>] Example: apicl(config-monitor-access) # source interface vpc pc1 leaf 101 102	Specifies the source interface vPC.
Step 13	drop enable Example: apicl(config-monitor-access-source) # drop enable	Enables the SPAN on drop feature, which captures all packets that are dropped in the ASIC and sends them to a pre-configured SPAN destination.
Step 14	[no] direction { rx tx both } Example: apicl(config-monitor-access-source) # direction tx	Specifies direction of traffic to be monitored. The direction can be configured independently for each source port range.
Step 15	[no] filter tenant <i>tenant-name</i> application <i>application-name</i> epg <i>epg-name</i> Example:	Filters traffic to be monitored. The filter can be configured independently for each source port range.

	Command or Action	Purpose
	<code>apic1(config-monitor-access-source)# filter tenant t1 application appl epg epg1</code>	
Step 16	exit Example: <code>apic1(config-monitor-access-source)# exit</code>	Returns to access monitor session configuration mode.
Step 17	[no] shutdown Example: <code>apic1(config-monitor-access)# no shut</code>	Disables (or enables) the monitoring session.

Examples

This example shows how to configure an ERSPAN access monitoring session.

```

apic1# configure terminal
apic1(config)# monitor access session mySession
apic1(config-monitor-access)# description "This is my access ERSPAN session"
apic1(config-monitor-access)# destination tenant t1 application appl epg epg1 destination-ip
 192.0.20.123 source-ip-prefix 10.0.20.1
apic1(config-monitor-access-dest)# erspan-id 100
apic1(config-monitor-access-dest)# ip dscp 42
apic1(config-monitor-access-dest)# ip ttl 16
apic1(config-monitor-access-dest)# mtu 9216
apic1(config-monitor-access-dest)# exit
apic1(config-monitor-access)# source interface eth 1/1 leaf 101
apic1(config-monitor-access-source)# direction tx
apic1(config-monitor-access-source)#drop enable
apic1(config-monitor-access-source)# filter tenant t1 application appl epg epg1
apic1(config-monitor-access-source)# exit
apic1(config-monitor-access)# no shut
apic1(config-monitor-access)# show run
# Command: show running-config monitor access session mySession
# Time: Fri Nov 6 23:55:35 2015
  monitor access session mySession
    description "This is my ERSPAN session"
    source interface eth 1/1 leaf 101
      direction tx
      filter tenant t1 application appl epg epg1
      exit
    destination tenant t1 application appl epg epg1 destination-ip 192.0.20.123
  source-ip-prefix 10.0.20.1
    ip dscp 42
    ip ttl 16
    erspan-id 9216
    mtu 9216
    exit
  exit

```

This example shows how to configure a port-channel as a monitoring source.

```
apicl(config-monitor-access)# source interface port-channel pc3 leaf 105
```

This example shows how to configure a one leg of a vPC as a monitoring source.

```
apicl(config-monitor-access)# source interface port-channel vpc3 leaf 105
```

This example shows how to configure a range of ports from FEX 101 as a monitoring source.

```
apicl(config-monitor-access)# source interface eth 101/1/1-2 leaf 105
```

Configuring ERSPAN in Fabric Mode Using the NX-OS-Style CLI

In the ACI fabric, a fabric mode ERSPAN configuration can be used for monitoring traffic originating from one or more fabric ports in leaf or spine nodes. Local SPAN is not supported in fabric mode.

For an ERSPAN session, the destination is always an endpoint group (EPG) which can be deployed anywhere in the fabric. The monitored traffic is forwarded to the destination wherever the EPG is moved. In the fabric mode, only fabric ports are allowed as source, but both leaf and spine switches are allowed.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: apicl# configure terminal	Enters global configuration mode.
Step 2	[no] monitor fabric session <i>session-name</i> Example: apicl(config)# monitor fabric session mySession	Creates a fabric monitoring session configuration.
Step 3	[no] description <i>text</i> Example: apicl(config-monitor-fabric)# description "This is my fabric ERSPAN session"	Adds a description for this monitoring session. If the text includes spaces, it must be enclosed in single quotes.
Step 4	[no] destination tenant <i>tenant-name</i> application <i>application-name</i> epg <i>epg-name</i> destination-ip <i>dest-ip-address</i> source-ip-prefix <i>src-ip-address</i> Example: apicl(config-monitor-fabric)# destination tenant t1 application app1 epg epg1 destination-ip 192.0.20.123 source-ip-prefix 10.0.20.1	Specifies the destination interface as a tenant and enters destination configuration mode.

	Command or Action	Purpose
Step 5	[no] erspan-id <i>flow-id</i> Example: apicl(config-monitor-fabric-dest)# erspan-id 100	Configures the ERSPAN ID for the ERSPAN session. The ERSPAN range is from 1 to 1023.
Step 6	[no] ip dscp <i>dscp-code</i> Example: apicl(config-monitor-fabric-dest)# ip dscp 42	Configures the differentiated services code point (DSCP) value of the packets in the ERSPAN traffic. The range is from 0 to 64.
Step 7	[no] ip ttl <i>ttl-value</i> Example: apicl(config-monitor-fabric-dest)# ip ttl 16	Configures the IP time-to-live (TTL) value for the ERSPAN traffic. The range is from 1 to 255.
Step 8	[no] mtu <i>mtu-value</i> Example: apicl(config-monitor-fabric-dest)# mtu 9216	Configures the maximum transmit unit (MTU) size for the ERSPAN session. The range is 64 to 9216 bytes.
Step 9	exit Example: apicl(config-monitor-fabric-dest)#	Returns to monitor access configuration mode.
Step 10	[no] source interface ethernet { <i>slot/port</i> <i>port-range</i> } switch <i>node-id</i> Example: apicl(config-monitor-fabric)# source interface eth 1/2 switch 101	Specifies the source interface port or port range.
Step 11	drop enable Example: apicl(config-monitor-fabric-source)# drop enable	Enables the SPAN on drop feature, which captures all packets that are dropped in the ASIC and sends them to a pre-configured SPAN destination.
Step 12	[no] direction { <i>rx</i> <i>tx</i> <i>both</i> } Example: apicl(config-monitor-fabric-source)# direction tx	Specifies direction of traffic to be monitored. The direction can be configured independently for each source port range.
Step 13	[no] filter tenant <i>tenant-name</i> bd <i>bd-name</i> Example: apicl(config-monitor-fabric-source)# filter tenant t1 bd bd1	Filters traffic by bridge domain.

	Command or Action	Purpose
Step 14	[no] filter tenant <i>tenant-name</i> vrf <i>vrf-name</i> Example: apicl(config-monitor-fabric-source) # filter tenant t1 vrf vrf1	Filters traffic by VRF.
Step 15	exit Example: apicl(config-monitor-fabric-source) # exit	Returns to access monitor session configuration mode.
Step 16	[no] shutdown Example: apicl(config-monitor-fabric) # no shut	Disables (or enables) the monitoring session.

Examples

This example shows how to configure an ERSPAN fabric monitoring session.

```

apicl# configure terminal
apicl(config)# monitor fabric session mySession
apicl(config-monitor-fabric)# description "This is my fabric ERSPAN session"
apicl(config-monitor-fabric)# destination tenant t1 application appl epg epg1 destination-ip
192.0.20.123 source-ip-prefix 10.0.20.1
apicl(config-monitor-fabric-dest)# erspan-id 100
apicl(config-monitor-fabric-dest)# ip dscp 42
apicl(config-monitor-fabric-dest)# ip ttl 16
apicl(config-monitor-fabric-dest)# mtu 9216
apicl(config-monitor-fabric-dest)# exit
apicl(config-monitor-fabric)# source interface eth 1/1 switch 101
apicl(config-monitor-fabric-source)# drop enable
apicl(config-monitor-fabric-source)# direction tx
apicl(config-monitor-fabric-source)# filter tenant t1 bd bd1
apicl(config-monitor-fabric-source)# filter tenant t1 vrf vrf1
apicl(config-monitor-fabric-source)# exit
apicl(config-monitor-fabric)# no shut

```

Configuring ERSPAN in Tenant Mode Using the NX-OS-Style CLI

In the ACI fabric, a tenant mode ERSPAN configuration can be used for monitoring traffic originating from endpoint groups within a tenant.

In the tenant mode, traffic originating from a source EPG is sent to a destination EPG within the same tenant. The monitoring of traffic is not impacted if the source or destination EPG is moved within the fabric.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: apic1# configure terminal	Enters global configuration mode.
Step 2	[no] monitor tenant <i>tenant-name</i> session <i>session-name</i> Example: apic1(config)# monitor tenant session mySession	Creates a tenant monitoring session configuration.
Step 3	[no] description <i>text</i> Example: apic1(config-monitor-tenant)# description "This is my tenant ERSPAN session"	Adds a description for this access monitoring session. If the text includes spaces, it must be enclosed in single quotes.
Step 4	[no] destination tenant <i>tenant-name</i> application <i>application-name</i> epg <i>epg-name</i> destination-ip <i>dest-ip-address</i> source-ip-prefix <i>src-ip-address</i> Example: apic1(config-monitor-tenant)# destination tenant t1 application appl epg epg1 destination-ip 192.0.20.123 source-ip-prefix 10.0.20.1	Specifies the destination interface as a tenant and enters destination configuration mode.
Step 5	[no] erspan-id <i>flow-id</i> Example: apic1(config-monitor-tenant-dest)# erspan-id 100	Configures the ERSPAN ID for the ERSPAN session. The ERSPAN range is from 1 to 1023.
Step 6	[no] ip dscp <i>dscp-code</i> Example: apic1(config-monitor-tenant-dest)# ip dscp 42	Configures the differentiated services code point (DSCP) value of the packets in the ERSPAN traffic. The range is from 0 to 64.
Step 7	[no] ip ttl <i>tll-value</i> Example: apic1(config-monitor-tenant-dest)# ip ttl 16	Configures the IP time-to-live (TTL) value for the ERSPAN traffic. The range is from 1 to 255.
Step 8	[no] mtu <i>mtu-value</i> Example: apic1(config-monitor-tenant-dest)# mtu 9216	Configures the maximum transmit unit (MTU) size for the ERSPAN session. The range is 64 to 9216 bytes.

	Command or Action	Purpose
Step 9	exit Example: apicl(config-monitor-tenant-dest)#	Returns to monitor access configuration mode.
Step 10	[no] source application <i>application-name</i> epg <i>epg-name</i> Example: apicl(config-monitor-tenant)# source application app2 epg epg5	Specifies the source interface port or port range.
Step 11	[no] direction {rx tx both} Example: apicl(config-monitor-tenant-source)# direction tx	Specifies direction of traffic to be monitored. The direction can be configured independently for each source port range.
Step 12	exit Example: apicl(config-monitor-tenant-source)# exit	Returns to access monitor session configuration mode.
Step 13	[no] shutdown Example: apicl(config-monitor-tenant)# no shut	Disables (or enables) the monitoring session.

Examples

This example shows how to configure an ERSPAN tenant monitoring session.

```

apicl# configure terminal
apicl(config)# monitor access session mySession
apicl(config-monitor-tenant)# description "This is my tenant ERSPAN session"
apicl(config-monitor-tenant)# destination tenant t1 application appl1 epg epg1 destination-ip
192.0.20.123 source-ip-prefix 10.0.20.1
apicl(config-monitor-tenant-dest)# erspan-id 100
apicl(config-monitor-tenant-dest)# ip dscp 42
apicl(config-monitor-tenant-dest)# ip ttl 16
apicl(config-monitor-tenant-dest)# mtu 9216
apicl(config-monitor-tenant-dest)# exit
apicl(config-monitor-tenant)# source application app2 epg epg5
apicl(config-monitor-tenant-source)# direction tx
apicl(config-monitor-tenant-source)# exit
apicl(config-monitor-tenant)# no shut

```

Configuring a Global SPAN-On-Drop Session Using the NX-OS-Style CLI

This section demonstrates how to create a global drop with all ports on a node as the SPAN source.

Procedure

Step 1 **configure terminal**

Enters global configuration mode.

Example:

```
apic1# configure terminal
```

Step 2 **[no] monitor fabric session *session-name***

Creates a fabric monitoring session configuration.

Example:

```
apic1(config)# monitor fabric session Spine301-GD-SOD
```

Step 3 **[no] description *text***

Adds a description for this monitoring session. If the text includes spaces, it must be enclosed in single quotes.

Example:

```
apic1(config-monitor-fabric)# description "This is my fabric ERSPAN session"
```

Step 4 **source global-drop switch**

Enables the SPAN on drop feature, which captures all packets that are dropped in the ASIC and sends them to a pre-configured SPAN destination.

Example:

```
apic1(config-monitor-fabric)# source global-drop switch
```

Step 5 **[no] destination tenant *tenant-name* application *application-name* epg *epg-name* destination-ip *dest-ip-address* source-ip-prefix *src-ip-address***

Specifies the destination interface as a tenant and enters destination configuration mode.

Example:

```
apic1(config-monitor-fabric-dest)# destination tenant ERSPAN application A1 epg E1
destination-ip 165.10.10.155 source-ip-prefix 22.22.22.22
```

Examples

This example shows how to configure a global SPAN-on-Drop session.

```
apic1# configure terminal
apic1(config)# monitor fabric session Spine301-GD-SOD
apic1(config-monitor-fabric)# source global-drop switch
apic1(config-monitor-fabric)# destination tenant ERSPAN application A1 epg E1 destination-ip
179.10.10.179 source-ip-prefix 31.31.31.31
```

Configuring SPAN Using the REST API

Configuring a Fabric Destination Group for an ERSPAN Destination Using the REST API

This section demonstrates how to use the REST API to configure a fabric destination group for an ERSPAN destination using the REST API. For a complete list of properties, see the *APIC Management Information Model Reference*.

Procedure

Configure a fabric destination group for an ERSPAN destination:

```
POST https://<APIC_IP>/api/node/mo/uni/infra.xml
<spanDestGrp annotation="" descr="" name="Dest2" nameAlias="" ownerKey="" ownerTag="">
  <spanDest annotation="" descr="" name="Dest2" nameAlias="" ownerKey="" ownerTag="">
    <spanRsDestEpg annotation="" dscp="unspecified" finalIp="0.0.0.0" flowId="1"
ip="179.10.10.179"
    mtu="1518"srcIpPrefix="20.20.20.2" tDn="uni/tn-ERSPAN/ap-A1/epg-E1" ttl="64" ver="ver2"
    verEnforced="no"/>
  </spanDest>
</spanDestGrp>
```

Configuring a Global Drop Source Group Using the REST API

This section demonstrates how to use the REST API to configure a global drop source group using the REST API. For a complete list of properties, see the *APIC Management Information Model Reference*.

Procedure

Configure a global drop source group:

```
POST https://<APIC_IP>/api/node/mo/uni/infra.xml
<spanSrcGrp adminSt="enabled" annotation="" descr="" name="Spine-402-GD-SOD" nameAlias="">
  <spanSrc annotation="" descr="" dir="both" name="402" nameAlias="" spanOnDrop="yes">
    <spanRsSrcToNode annotation="" tDn="topology/pod-1/node-402"/>
  </spanSrc><spanSpanLbl annotation="" descr="" name="402-dst-179" nameAlias=""
tag="yellow-green"/>
</spanSrcGrp>
```

Configuring a Leaf Port as a SPAN Destination Using the REST API

This section demonstrates how to use the REST API to configure a leaf port as a SPAN destination using the REST API. For a complete list of properties, see the *APIC Management Information Model Reference*.

Procedure

Configure a leaf port as a SPAN destination:

```
POST https://<APIC_IP>/api/node/mo/uni/infra.xml
<spanDestGrp annotation="" descr="" name="Dest4" nameAlias="" ownerKey="" ownerTag="">
  <spanDest annotation="" descr="" name="Dest4" nameAlias="" ownerKey="" ownerTag="">
    <spanRsDestPathEp annotation="" mtu="1518"
    tDn="topology/pod-1/paths-301/pathep-[eth1/18]"/>
  </spanDest>
</spanDestGrp>
```

Configuring a SPAN Access Source Group Using the REST API

This section demonstrates how to use the REST API to configure a SPAN access source group using the REST API. For a complete list of properties, see the *APIC Management Information Model Reference*.

Procedure

Configure a SPAN access source group:

```
POST https://<APIC_IP>/api/node/mo/uni/infra.xml
<spanSrcGrp adminSt="enabled" annotation="" descr="" name="Test-Src2" nameAlias=""
ownerKey=""
ownerTag="">
  <spanSrc annotation="" descr="" dir="both" name="Src1" nameAlias="" ownerKey=""
ownerTag=""
spanOnDrop="yes">
  <spanRsSrcToPathEp annotation="" tDn="topology/pod-1/paths-301/pathep-[eth1/1]"/>
</spanSrc>
  <spanSpanLbl annotation="" descr="" name="Dest1" nameAlias="" ownerKey="" ownerTag=""
tag="yellow-green"/>
</spanSrcGrp>
```

Configuring a SPAN Fabric Source Group Using the REST API

This section demonstrates how to use the REST API to configure a SPAN fabric source group using the REST API. For a complete list of properties, see the *APIC Management Information Model Reference*.

Procedure

Configure a SPAN fabric source group:

```
POST https://<APIC_IP>/api/node/mo/uni/infra.xml
<spanSrcGrp adminSt="enabled" annotation="" descr="" name="Test-Src2" nameAlias="" ownerKey=""
```

```

ownerTag="">
  <spanSrc annotation="" descr="" dir="both" name="Src1" nameAlias="" ownerKey=""
  ownerTag="" spanOnDrop="yes">
  <spanRsSrcToPathEp annotation="" tDn="topology/pod-1/paths-301/pathep-[eth1/51]"/>
  </spanSrc>
  <spanSpanLbl annotation="" descr="" name="Dest2" nameAlias="" ownerKey="" ownerTag=""
  tag="yellow-green"/>
</spanSrcGrp>

```

Configuring an Access Destination Group for an ERSPAN Destination Using the REST API

This section demonstrates how to use the REST API to configure an access destination group for an ERSPAN destination using the REST API. For a complete list of properties, see the *APIC Management Information Model Reference*.

Procedure

Configure an access destination group for an ERSPAN destination.

```

POST https://<APIC_IP>/api/node/mo/uni/infra.xml
<spanDestGrp annotation="" descr="" name="Dest4" nameAlias="" ownerKey=""
  ownerTag="">
  <spanDest annotation="" descr="" name="Dest4" nameAlias="" ownerKey=""
  ownerTag="">
  <spanRsDestPathEp annotation="" mtu="1518" tDn="topology/pod-1/paths-301/pathep-
  [eth1/18]"/>
  </spanDest>
</spanDestGrp>

```

Configuring a SPAN Filter With Extended Filters Using the REST API

The following example shows you how configure a SPAN Filter using the REST API.

Procedure

To configure a SPAN Filter using the Rest API:

Example:

```

URL: {{apic-host}}/api/node/mo/.xml
BODY:
<polUni>
  <infraInfra dn="uni/infra">
    <spanSrcGrp adminSt="enabled" descr="" dn="uni/infra/srcgrp-local1" nameAlias=""
  ownerKey=""
  ownerTag="">
      <spanRsSrcGrpToFilterGrp tDn="uni/infra/filtergrp-two" />
      <spanSrc descr="" dir="both" name="src1" nameAlias="" ownerKey="" ownerTag="">
      <spanRsSrcToPathEp tDn="topology/pod-1/paths-101/pathep-[eth1/15]" />
      </spanSrc>
      <spanSpanLbl descr="" name="dest1" nameAlias="" ownerKey="" ownerTag="" tag=
      "yellow-green" />
    </spanSrcGrp>
  </infraInfra>
</polUni>

```

```

        </spanSrcGrp>
        <spanDestGrp annotation="" descr="" dn="uni/infra/destgrp-dest1" nameAlias=""
ownerKey=""
        ownerTag="">
        <spanDest annotation="" descr="" name="destg" nameAlias="" ownerKey="" ownerTag="">

            <spanRsDestPathEp annotation="" mtu="1518" tDn="topology/pod-1/paths-101/pathep-
[eth1/7]" />
            </spanDest>
        </spanDestGrp>
        <spanFilterGrp name="two">
        <spanFilterEntry name="udp_two" ipProto="udp" srcAddr="1002::1/64"
dstAddr="1001::1/64"
        srcPortFrom="1" srcPortTo="2" dstPortFrom="1" dstPortTo="2">
            <spanExtendedFltEntry name="arun1" dscpFrom="0" dscpTo="10" dot1pFrom="0"
dot1pTo="7"
            tcpFlags="128" v6FlowLabel="1522" forwardingVal="switched" />
            </spanFilterEntry>
        </spanFilterGrp>
    </infraInfra>
</polUni>

```

Using Statistics

Statistics provide real-time measures of observed object and enable trend analysis and troubleshooting. Statistics gathering can be configured for ongoing or on-demand collection and can be collected in cumulative counters and gauges.

Policies define what statistics are gathered, at what intervals, and what actions to take. For example, a policy could raise a fault on an EPG if a threshold of dropped packets on an ingress VLAN is greater than 1000 per second.

Statistics data are gathered from a variety of sources, including interfaces, VLANs, EPGs, application profiles, ACL rules, tenants, or internal APIC processes. Statistics accumulate data in 5-minute, 15-minute, 1-hour, 1-day, 1-week, 1-month, 1-quarter, or 1-year sampling intervals. Shorter duration intervals feed longer intervals. A variety of statistics properties are available, including last value, cumulative, periodic, rate of change, trend, maximum, min, average. Collection and retention times are configurable. Policies can specify if the statistics are to be gathered from the current state of the system or to be accumulated historically or both. For example, a policy could specify that historical statistics be gathered for 5-minute intervals over a period of 1 hour. The 1 hour is a moving window. Once an hour has elapsed, the incoming 5 minutes of statistics are added, and the earliest 5 minutes of data are abandoned.



Note The maximum number of 5-minute granularity sample records is limited to 12 samples (one hour of statistics). All other sample intervals are limited to 1,000 sample records. For example, hourly granularity statistics can be maintained for up to 41 days.

Viewing Statistics in the GUI

You can view statistics for many objects using the APIC GUI, including application profiles, physical interfaces, bridge domains, and fabric nodes. To view statistics in the GUI, choose the object in the **navigation** pane and click the STATS tab.

Follow these steps to view statistics for an interface:

Procedure

-
- Step 1** On the menu bar, choose **FABRIC > INVENTORY**.
 - Step 2** In the **Navigation** pane, choose a pod.
 - Step 3** Expand the pod, and expand a switch.
 - Step 4** In the **Navigation** pane, expand **Interfaces** and choose **eth1/1**.
 - Step 5** In the **Work** pane, choose the **STATS** tab.
-

The APIC displays interface statistics.

Example

What to do next

You can use the following icons in the **Work** pane to manage how the APIC displays statistics:

- Refresh—Manually refreshes statistics.
- Show Table View—Toggles between table and chart views.
- Start or Stop Stats—Enables or disables automatic refresh for statistics.
- Select Stats—Specifies the counters and sample interval to display.
- Download Object as XML—Downloads the object in XML format.
- Measurement Type (Gear icon)—Specifies the statistics measurement type. Options include cumulative, periodic, average, or trend.

Switch Statistics Commands

You can use the following commands to display statistics on ACI leaf switches.

Command	Purpose
Legacy Cisco Nexus show/clear commands	For more information, see <i>Cisco Nexus 9000 Series NX-OS Configuration Guides</i> .

Command	Purpose
show platform internal counters port [<i>port_num</i> detail nz { internal [<i>nz</i> <i>int_port_num</i>]}]	<p>Displays spine port statistics</p> <ul style="list-style-type: none"> • <i>port_num</i>—Front port number without the slot. • detail—Returns SNMP, class and forwarding statistics. • nz—Displays only non-zero values. • internal—Displays internal port statistics. • <i>int_port_num</i>—Internal logical port number. For example, for BCM-0/97, enter 97. <p>Note If there is a link reset, the counters will be zeroed out on the switch. The conditions of counter reset include the following:</p> <ul style="list-style-type: none"> • accidental link reset • manually enabled port (after port is disabled)
show platform internal counters vlan [<i>hw_vlan_id</i>]	Displays VLAN statistics.
show platform internal counters tep [<i>tunnel_id</i>]	Displays TEP statistics.
show platform internal counters flow [<i>rule_id</i> { dump [<i>asic inst</i>] [slice direction index hw_index]}]	Displays flow statistics.
clear platform internal counters port [<i>port_num</i> { internal [<i>int_port_num</i>]}]	Clears port statistics.
clear platform internal counters vlan [<i>hw_vlan_id</i>]	Clears VLAN counters.
debug platform internal stats logging level <i>log_level</i>	Sets the debug logging level.
debug platform internal stats logging { err trace flow }	Sets the debug logging type.

Managing Statistics Thresholds Using the GUI

Procedure

-
- Step 1** On the menu bar, choose **Fabric > Fabric Policies**.
 - Step 2** In the **Navigation** pane, click + to expand **Monitoring Policies**.
 - Step 3** In the **Navigation** pane, expand the monitoring policy name (such as Default).
 - Step 4** Click **Stats Collection Policies**.
 - Step 5** In the **Stats Collection Policies** window, choose a **Monitoring Object** and **Stats Type** for which to set a threshold value..

- Step 6** In the **Work** pane, Click the + icon below **CONFIG THRESHOLDS**.
- Step 7** In the **THRESHOLDS FOR COLLECTION** window, click + to add a threshold.
- Step 8** In the **Choose a Property** window, choose a statistics type.
- Step 9** In the **EDIT STATS THRESHOLD** window, specify the following threshold values:
- Normal Value—A valid value of the counter.
 - Threshold Direction—Indicates whether the threshold is a maximum or minimum value.
 - Rising Thresholds (Critical, Major, Minor, Warning)—Triggered when the value exceeds the threshold.
 - Falling Thresholds (Critical, Major, Minor, Warning)—Triggered when the value drops below the threshold.
- Step 10** You can specify a set and reset value for rising and falling thresholds. The set value specifies when a fault is triggered; the reset value specifies when the fault is cleared.
- Step 11** Click **SUBMIT** to save the threshold value.
- Step 12** In the **THRESHOLDS FOR COLLECTION** window, click **CLOSE**.

Statistics Troubleshooting Scenarios

The following table summarizes common statistics troubleshooting scenarios for the Cisco APIC.

Problem	Solution
The APIC does not enforce a configured monitoring policy	<p>The problem occurs when a monitoring policy is in place but the APIC does not perform a corresponding action, such as collecting the statistics or acting on a trigger threshold. Follow these steps to resolve the issue:</p> <ul style="list-style-type: none"> • Verify that monPolDn points to the correct monitoring policy. • Ensure that the selectors are configured correctly and that there are no faults. • For Tenant objects, check the relation to the monitoring policy.
Some configured statistics are missing.	<p>Follow these steps to resolve the issue:</p> <ul style="list-style-type: none"> • Review the statistics that are disabled by default within the monitoring policy and collection policy. • Review the collection policy to determine if the statistics are disabled by default or disabled for certain intervals. • Review the statistics policy to determine if the statistics are disabled by default or disabled for certain intervals. <p>Note Except for fabric health statistics, 5 minute statistics are stored on the switch and are lost when the switch reboots.</p>

Problem	Solution
Statistics or history are not maintained for the configured time period.	<p>Follow these steps to resolve the issue:</p> <ul style="list-style-type: none"> • Review the collection settings; if configured at the top level of the monitoring policy, the statistics can be overridden for a specific object or statistics type. • Review the collection policy assigned to the monitoring object. Confirm that the policy is present and review the administrative state, and history retention values. • Verify that the statistics type is configured correctly.
Some statistics are not maintained for the full configured interval.	<p>Review whether the configuration exceeds the maximum historical record size. The limitations are as follows:</p> <ul style="list-style-type: none"> • Switch statistics for 5 minute granularity are limited to 12 samples (1 hour of 5 minute granular statistics). • There is a hard limit of 1000 samples. For example, hourly granular statistics can be maintained for up to 41 days.
An export policy is configured but the APIC does not export statistics.	<p>Follow these steps to resolve the issue:</p> <ul style="list-style-type: none"> • Check the status object for the destination policy. • On the node that is expected to export the statistics check the export status object and look at the export status and details properties. Aggregated EPG stats are exported every 15 minutes from APIC nodes. Other statistics are exported from source nodes every 5 minutes. For example, if an EPG is deployed to two leaf switches and configured to export EPG aggregation parts, then those parts are exported from the nodes every 5 minutes. • Review whether the configuration exceeds the maximum number of export policies. The maximum number of statistics export policies is approximately equal to the number of tenants. <p>Note Each tenant can have multiple statistics export policies and multiple tenants can share the same export policy, but the total number number of policies is limited to approximately the number of tenants.</p>
5 Minute Statistics Fluctuate	The APIC system reports statistics every 5 minutes, sampled approximately every 10 seconds. The number of samples taken in 5 minutes may vary, because there are slight time variances when the data is collected. As a result, the statistics might represent a slightly longer or shorter time period. This is expected behavior.
Some historical statistics are missing.	For more information, see Statistics Cleanup .

Statistics Cleanup

The APIC and switches clean up statistics as follows:

- Switch—The switch cleans up statistics as follows:

- 5 minute statistics on switches are purged if no counter value is reported for 5 minutes. This situation can occur when an object is deleted or statistics are disabled by a policy.
- Statistics of larger granularity are purged if statistics are missing for more than one hour, which can occur when:
 - Statistics are disabled by a policy.
 - A switch is disconnected from an APIC for more than one hour.
- The switch cleans up statistics for deleted objects after 5 minutes. If an object is recreated within this time, statistics counts remain unchanged.
- Disabled object statistics are deleted after 5 minutes.
- If the system state changes so that statistics reporting is disabled for 5 minutes, this switch cleans up statistics.
- APIC—The APIC cleans up objects including interfaces, EPGs, temperature sensors, and health statistics after one hour.

Using Syslog

About Syslog

During operation, a fault or event in the Cisco Application Centric Infrastructure (ACI) system can trigger the sending of a system log (syslog) message to the console, to a local file, and to a logging server on another system. A system log message typically contains a subset of information about the fault or event. A system log message can also contain audit log and session log entries.



Note For a list of syslog messages that the APIC and the fabric nodes can generate, see http://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/syslog/guide/aci_syslog/ACI_SysMsg.html.

Many system log messages are specific to the action that a user is performing or the object that a user is configuring or administering. These messages can be the following:

- Informational messages, providing assistance and tips about the action being performed
- Warning messages, providing information about system errors related to an object, such as a user account or service profile, that the user is configuring or administering

In order to receive and monitor system log messages, you must specify a syslog destination, which can be the console, a local file, or one or more remote hosts running a syslog server. In addition, you can specify the minimum severity level of messages to be displayed on the console or captured by the file or host. The local file for receiving syslog messages is `/var/log/external/messages`.

A syslog source can be any object for which an object monitoring policy can be applied. You can specify the minimum severity level of messages to be sent, the items to be included in the syslog messages, and the syslog destination.

You can change the display format for the Syslogs to NX-OS style format.

Additional details about the faults or events that generate these system messages are described in the *Cisco APIC Faults, Events, and System Messages Management Guide*, and system log messages are listed in the *Cisco ACI System Messages Reference Guide*.



Note Not all system log messages indicate problems with your system. Some messages are purely informational, while others may help diagnose problems with communications lines, internal hardware, or the system software.

Creating a Syslog Destination and Destination Group

This procedure configures syslog data destinations for logging and evaluation. You can export syslog data to the console, to a local file, or to one or more syslog servers in a destination group.

Procedure

-
- Step 1** In the menu bar, click **Admin**.
- Step 2** In the submenu bar, click **External Data Collectors**.
- Step 3** In the **Navigation** pane, expand **Monitoring Destinations**.
- Step 4** Right-click **Syslog** and choose **Create Syslog Monitoring Destination Group**.
- Step 5** In the **Create Syslog Monitoring Destination Group** dialog box, perform the following actions:
- In the group and profile **Name** field, enter a name for the monitoring destination group and profile.
 - In the group and profile **Format** field, choose the format for Syslog messages.

The default is **aci**, or the RFC 5424 compliant message format, but you can choose to set it to the NX-OS style format instead.
 - In the group and profile **Admin State** drop-down list, choose **enabled**.
 - To enable sending of syslog messages to a local file, choose **enabled** from the Local File Destination **Admin State** drop-down list and choose a minimum severity from the Local File Destination **Severity** drop-down list.

The local file for receiving syslog messages is `/var/log/external/messages`.
 - To enable sending of syslog messages to the console, choose **enabled** from the Console Destination **Admin State** drop-down list and choose a minimum severity from the Console Destination **Severity** drop-down list.
 - Click **Next**.
 - In the **Create Remote Destinations** area, click + to add a remote destination.
- Caution** There is a risk of hostname resolution failure for remote syslog destinations if the DNS server that you specify is configured to be reachable over in-band connectivity. To avoid the issue, configure the syslog server using the IP address, or if you use a hostname, ensure that the DNS server is reachable over an out-of-band interface.
- Step 6** In the **Create Syslog Remote Destination** dialog box, perform the following actions:
- In the **Host** field, enter an IP address or a fully qualified domain name for the destination host.
 - (Optional) In the **Name** field, enter a name for the destination host.
 - In the **Admin State** field, click the **enabled** radio button.

- d) (Optional) Choose a minimum severity **Severity**, a **Port** number, and a syslog **Facility**.
The **Facility** is a number that you can optionally use to indicate which process generated the message, and can then be used to determine how the message will be handled at the receiving end.
- e) In the 5.2(3) release and later, in the **Transport** field, choose the transport protocol to use for the messages.
- For releases prior to release 5.2(4), choose either **tcp** or **udp** as the transport protocol to use for the messages.
 - In the 5.2(4) release and later, **ssl** is also an option for the transport protocol to use for the messages. This feature enables a Cisco ACI switch (acting as a client) to make a secure, encrypted outbound connection to remote syslog servers (acting as a server) supporting secure connectivity for logging. With authentication and encryption, this feature allows for a secure communication over an insecure network.
- Note that you must also upload the necessary SSL certificate if you select **ssl** as the transport protocol to use for the messages. You can upload the necessary SSL certificate by navigating to the **Create Certificate Authority** window:
- Admin > AAA > Security > Public Key Management > Certificate Authorities, then Actions > Create Certificate Authority**
- The default option for the transport protocol is **udp**.
- f) From the **Management EPG** drop-down list, choose the management endpoint group.
- g) Click **OK**.

Step 7 (Optional) To add more remote destinations to the remote destination group, click + again and repeat the steps in the **Create Syslog Remote Destination** dialog box

Step 8 Click **Finish**.

Creating a Syslog Source

A syslog source can be any object for which an object monitoring policy can be applied.

Before you begin

Create a syslog monitoring destination group.

Procedure

-
- Step 1** From the menu bar and the navigation frame, navigate to a **Monitoring Policies** menu for the area of interest. You can configure monitoring policies for tenants, fabric, and access.
- Step 2** Expand **Monitoring Policies**, then select and expand a monitoring policy.
- Under **Fabric > Fabric Policies > Monitoring Policies > Common Policy** is a basic monitoring policy that applies to all faults and events and is automatically deployed to all nodes and controllers in the fabric. Alternatively, you can specify an existing policy with a more limited scope.
- Step 3** Under the monitoring policy, click **Callhome/SNMP/Syslog**.

- Step 4** In the **Work** pane, choose **Syslog** from the **Source Type** drop-down list.
- Step 5** From the **Monitoring Object** list, choose a managed object to be monitored.
- If the desired object does not appear in the list, follow these steps:
- Click the Edit icon to the right of the **Monitoring Object** drop-down list.
 - From the **Select Monitoring Package** drop-down list, choose an object class package.
 - Select the checkbox for each object that you want to monitor.
 - Click **Submit**.
- Step 6** In a tenant monitoring policy, if you select a specific object instead of **All**, a **Scope** selection appears.
- In the **Scope** field, select a radio button to specify the system log messages to send for this object:
- **all**—Send all events and faults related to this object
 - **specific event**—Send only the specified event related to this object. From the **Event** drop-down list, choose the event policy.
 - **specific fault**—Send only the specified fault related to this object. From the **Fault** drop-down list, choose the fault policy.
- Step 7** Click + to create a syslog source.
- Step 8** In the **Create Syslog Source** dialog box, perform the following actions:
- In the **Name** field, enter a name for the syslog source.
 - From the **Min Severity** drop-down list, choose the minimum severity of system log messages to be sent.
 - In the **Include** field, check the checkboxes for the type of messages to be sent.
 - From the **Dest Group** drop-down list, choose the syslog destination group to which the system log messages will be sent.
 - Click **Submit**.
- Step 9** (Optional) To add more syslog sources, click + again and repeat the steps in the **Create Syslog Source** dialog box
-

Using Traceroute

About Traceroute

The traceroute tool is used to discover the routes that packets actually take when traveling to their destination. Traceroute identifies the path taken on a hop-by-hop basis and includes a time stamp at each hop in both directions. You can use traceroute to test the connectivity of ports along the path between the generating device and the device closest to the destination. If the destination cannot be reached, the path discovery traces the path up to the point of failure.

A traceroute that is initiated from the tenant endpoints shows the default gateway as an intermediate hop that appears at the ingress leaf switch.

Traceroute supports a variety of modes, including:

- Endpoint-to-endpoint, and leaf-to-leaf (tunnel endpoint, or TEP to TEP)
- Endpoint-to-external-IP

- External-IP-to-endpoint
- External-IP-to-external-IP

Traceroute discovers all paths across the fabric, discovers point of exits for external endpoints, and helps to detect if any path is blocked.

Traceroute Guidelines and Restrictions

- When the traceroute source or destination is an endpoint, the endpoint must be dynamic and not static. Unlike a dynamic endpoint (fv:CEp), a static endpoint (fv:StCEp) does not have a child object (fv:RsCEpToPathEp) that is required for traceroute.
- Traceroute works for IPv6 source and destinations but configuring source and destination IP addresses across IPv4 and IPv6 addresses is not allowed.
- See the *Verified Scalability Guide for Cisco ACI* document for traceroute-related limits.
- When an endpoint moves from one ToR switch to a different ToR switch that has a new MAC address (one that is different than the MAC address that you specified while configuring the traceroute policy), the traceroute policy shows "missing-target" for the endpoint. In this scenario you must configure a new traceroute policy with the new MAC address.
- When performing a traceroute for a flow involving the policy-based redirect feature, the IP address used by the leaf switch to source the time-to-live (TTL) expired message when the packet goes from the service device to the leaf switch may not always be the IP address of the bridge domain's switch virtual interface (SVI) of the service device. This behavior is cosmetic and does not indicate that the traffic is not taking the expected path.

Performing a Traceroute Between Endpoints

Procedure

-
- Step 1** In the menu bar, click **Tenants**.
- Step 2** In the submenu bar, click the tenant that contains the source endpoint.
- Step 3** In the **Navigation** pane, expand the tenant and expand **Policies > Troubleshoot**.
- Step 4** Under **Troubleshoot**, right-click on one of the following traceroute policies:
- **Endpoint-to-Endpoint Traceroute Policies** and choose **Create Endpoint-to-Endpoint Traceroute Policy**
 - **Endpoint-to-External-IP Traceroute Policies** and choose **Create Endpoint-to-External-IP Traceroute Policy**
 - **External-IP-to-Endpoint Traceroute Policies** and choose **Create External-IP-to-Endpoint Traceroute Policy**
 - **External-IP-to-External-IP Traceroute Policies** and choose **Create External-IP-to-External-IP Traceroute Policy**
- Step 5** Enter the appropriate values in the dialog box fields and click **Submit**.

Note For the description of a field, click the help icon (?) in the top-right corner of the dialog box.

Step 6 In the **Navigation** pane or the **Traceroute Policies** table, click the traceroute policy. The traceroute policy is displayed in the **Work** pane.

Step 7 In the **Work** pane, click the **Operational** tab, click the **Source Endpoints** tab, and click the **Results** tab.

Step 8 In the **Traceroute Results** table, verify the path or paths that were used in the trace.

Note

- More than one path might have been traversed from the source node to the destination node.
- For readability, you can increase the width of one or more columns, such as the **Name** column.

Using the Troubleshooting Wizard

The Troubleshooting Wizard allows you to understand and visualize how your network is behaving, which can ease your networking concerns should issues arise. For example, you might have two endpoints that are having intermittent packet loss, but you do not understand why. Using the Troubleshooting Wizard, you can evaluate the issue so that you can effectively resolve the issue rather than logging onto each machine that you suspect to be causing this faulty behavior.



This wizard allows you (the administrative user) to troubleshoot issues that occur during specific time frames for the chosen source and destination. You can define a time window in which you want to perform the debug, and you can generate a troubleshooting report that you can send to TAC.

Getting Started with the Troubleshooting Wizard

Before you start using the Troubleshooting Wizard, you must be logged on as an Administrative user. Then, you must designate a source and destination and choose a time window for your troubleshooting session. The time window is used for retrieving events, fault records, deployment records, audit logs, and statistics.

As you navigate through the screens of the Troubleshooting Wizard, you have the option to take a screen shot

at any time and send it to a printer or save it as a PDF by clicking the **Print** icon () at the top, right side

of the screen. There are also Zoom In and Zoom Out icons ( ) that you can use to modify your view of any screen.

**Note**

- You cannot modify the source and destination after you click either **Generate Report** or **Submit**. If you want to change the source and destination information after you have entered it, you must delete the current session and start a new session.
- You cannot modify the description and time window on the first page of the wizard after you click **Submit**.
- You cannot use static IP address endpoints with the Troubleshooting Wizard.
- Any endpoints that you specify must be under an EPG.
- If you specify either the source or destination to be an external IP address, add the external subnets in the L3Out External EPG configuration. For information about configuring external subnets for external EPGs, see the "Routed Connectivity to External Networks" chapter of the *Cisco APIC Layer 3 Networking Configuration Guide*.

To set up your troubleshooting session information:

Procedure

Step 1 Choose **Operations > Visibility & Troubleshooting**.

The **Visibility & Troubleshooting** screen appears.

Step 2 In the **Session Name** field, choose an existing troubleshooting session using the drop-down list or create a new session by entering a name.

Step 3 In the **Session Type** drop-down list, choose the desired session type.

- **Endpoint to Endpoint:** The source and destination are both internal endpoints.

You should choose source and destination endpoints from the same tenant, or some of the troubleshooting functionality may be impacted, as explained later in this document. With this session type, you cannot use atomic counters when both endpoints connect to the same set of leaf switches.

- **Endpoint to External IP:** The source is an internal endpoint, while the destination is an external IP address.
- **External IP to Endpoint:** The source is an external IP address, while the destination is an internal endpoint.
- **External IP to External IP:** The source and destination are both external IP addresses. You can choose this type starting in the 3.2(6) release. With this session type, you cannot use traceroute, atomic counters, nor latency.

Step 4 (Optional) Enter a description in the **Description** field to provide additional information.

Step 5 Enter the source information in the **Source** area.

- If you chose a session type of **Endpoint to Endpoint** or **Endpoint to External IP**, enter a MAC, IPv4, or IPv6 address, or a VM name, then click **Search**.

You can enter a MAC address only if the session type is **Endpoint to Endpoint** and both endpoints' MAC address do not have any IP addresses learned from them.

A box appears that displays one or more rows with detailed information to help you make a selection. Each row shows that the IP address (in the **IP** column) you entered is in a specific endpoint group (in the **EPG** column), which belongs to a certain application (in the **Application** column), which is in a particular tenant (in the **Tenant** column). The leaf switch number, FEX number, and port details are shown in the **Learned At** column.

- If you chose a session type of **External IP to Endpoint**, enter the external IP address.
- If you chose a session type of **External IP to External IP**, enter the external IP address and distinguished name of the external Layer 3 outside network.

Step 6 Enter the destination information in the **Destination** area.

- If you chose a session type of **Endpoint to Endpoint** or **External IP to Endpoint**, enter a MAC, IPv4, or IPv6 address, or a VM name, then click **Search**.

You can enter a MAC address only if the session type is **Endpoint to Endpoint** and both endpoints' MAC address do not have any IP addresses learned from them.

A box appears that displays one or more rows with detailed information to help you make a selection. Each row shows that the IP address (in the **IP** column) you entered is in a specific endpoint group (in the **EPG** column), which belongs to a certain application (in the **Application** column), which is in a particular tenant (in the **Tenant** column). The leaf switch number, FEX number, and port details are shown in the **Learned At** column.

- If you chose a session type of **Endpoint to External IP**, enter the external IP address.
- If you chose a session type of **External IP to External IP**, enter the external IP address and distinguished name of the external Layer 3 outside network.

Step 7 In the **Time Window** area, specify a time window.

The **Time Window** is used for debugging an issue that occurred during a specific time frame in the past, and is used for retrieving events, all records, deployment records, audit logs, and statistics. There are two sets of windows: one for all records and one for individual leaf switches (or nodes).

By default, you can specify a rolling time window based on any number of minutes that you specify in the **Latest Minutes** field. The default is 240 minutes. The session contains data for the past number of minutes that you specify that precede the time that you created the session.

If you put a check in the **Use fixed time** box, you can specify a fixed time window for the session in the **From** and **To** fields. The session contains data starting from the **From** time to the **To** time.

Step 8 Click **Submit** to begin your troubleshooting session.

After a short delay, the topology diagram for your troubleshooting session appears.

Generating Troubleshooting Reports

You can generate a troubleshooting report in several formats, including JSON, XML, PDF, and HTML. Once you select a format, you can download the report (or schedule a download of the report) and use it for offline analysis or you can send it to TAC so that a support case can be created.

To generate a troubleshooting report:

Procedure

-
- Step 1** From the bottom right corner of the screen, click **GENERATE REPORT**.
The **Generate Report** dialog box appears.
- Step 2** Choose an output format from the Report Format drop-down menu (**XML**, **HTML**, **JSON**, or **PDF**).
- Step 3** If you want to schedule the download of the report to happen immediately, click the **Now > SUBMIT**.
An **Information** box appears indicating where to obtain the report once it has been generated.
- Step 4** To schedule the generation of the report for a later time, choose a schedule by clicking **Use a scheduler > Scheduler** drop-down menu then choose either an existing schedule or create a new one by clicking **Create Scheduler**.
The **CREATE TRIGGER SCHEDULE** dialog appears.
- Step 5** Enter information for the **Name**, **Description** (optional), and **Schedule Windows** fields.
- Note** For more information on how to use the **SCHEDULER**, please refer to the online help.
- Step 6** Click **SUBMIT**.
-

The reports take some time to generate (from a couple of minutes to up to ten minutes), depending on the size of the fabric and how many faults or events exist. A status message displays while the report is being generated. To retrieve and view the troubleshooting report, click **SHOW GENERATED REPORTS**.

Supply the credentials (**User Name** and **Password**) of the server in the **Authentication Required** window. The troubleshooting report is then downloaded locally to your system.

The **ALL REPORTS** window appears showing a list of all the reports that have been generated, including the one you just triggered. From there, you can click the link to either download or immediately view the report, depending on the output file format you chose (for example, if the file is a PDF, it may open immediately in your browser).


Topology in the Troubleshooting Wizard

This section explains the topology in the Troubleshooting Wizard. The topology shows how the source and destination are connected to the fabric, what the network path is from the source to the destination, and what the intermediate switches are.

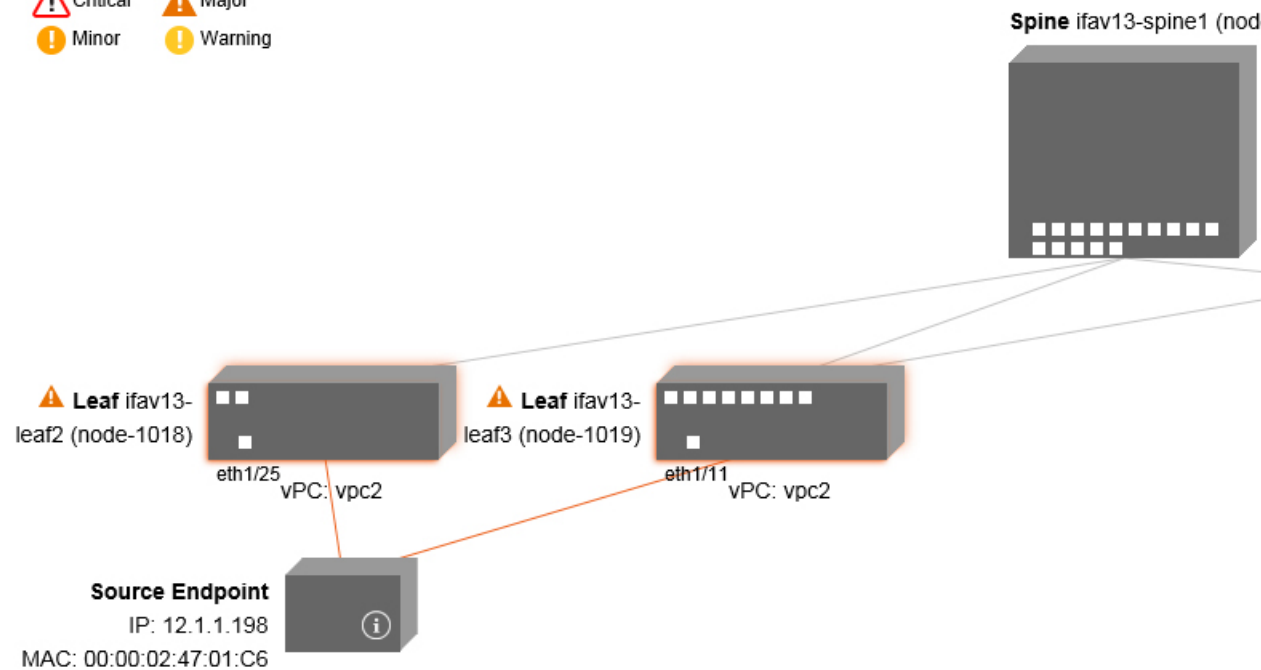
The source is displayed on the left side of the topology and the destination is on the right, as shown in the following wizard topology diagram.



Note This wizard topology only shows the leaf switches, spine switches, and FEXes of the devices involved in the traffic from the source to the destination. However, there may be many other leaf switches (tens or hundreds of leaf switches and many other spine switches) that exist.

This topology also shows links, ports, and devices. If you hover over the  icon), you can see the tenant that the source or destination belongs to, which application it belongs to, and the traffic encapsulation it is using (such as VLAN).

There is a color legend on the left side of the screen (shown as follows) that describes the severity levels associated with each color in the topology diagram (for example, critical versus minor).



Hovering over items such as boxes or ports in the topology provides more detailed information. If the port or link has a color, this means that there is a problem for you to troubleshoot. For example, if the color is red or orange, this indicates that there is a fault on a port or link. If the color is white, then there are no faults that exist. If the link has a number in a circle, it indicates how many parallel links between the same two nodes are affected by a fault of the severity given by the color of the circle. Hovering over a port allows you to see which port is connected to the source.

Right-clicking on a leaf switch allows you to access the console of the switch. A pop-up window appears that allows you to log into that device.




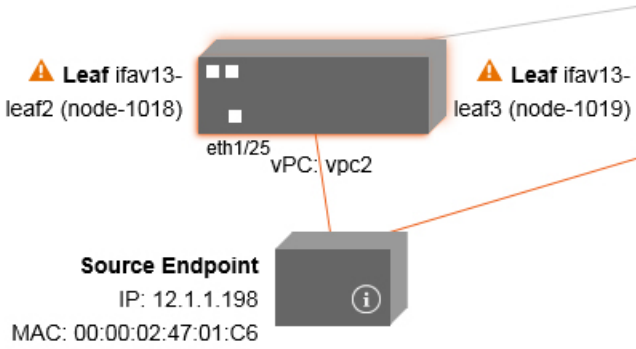
Note

- If there are Layer 4 to Layer 7 services (firewall and load balancer), they will be shown in the topology as well.
- For a topology with the load balancer, the destination is expected to be the virtual IP (VIP) address.
- When the source or destination is behind an ESX server, the ESX is shown in the topology.

Using the Faults Troubleshooting Screen

This procedure describes how to use the faults of the Troubleshooting Wizard.

Procedure

	Command or Action	Purpose
Step 1	Click Faults in the Navigation pane to begin using the Faults troubleshooting screen.	<p>The Faults screen shows the topology that connects the source and destination that you previously selected as well as the faults that were found. Only faults for the designated communication are shown. Wherever there are faults, they are highlighted in a certain color to convey the severity. Refer to the color legend at the top of the screen to understand the severity levels associated with each color. White boxes indicate that there are no issues to troubleshoot in that particular area.</p> <p>This topology also shows the relevant leaf switches, spine switches, and FEXes for your troubleshooting session. Hovering over items, such as leaf switches, spine switches, and FEXes, or clicking on faults provides more detailed information for analysis.</p> <p>  </p> 
Step 2	Click on a fault to display a dialog box with the Drop Stats , Contract Drops , and Traffic Stats tabs that contain more detailed information for analysis.	

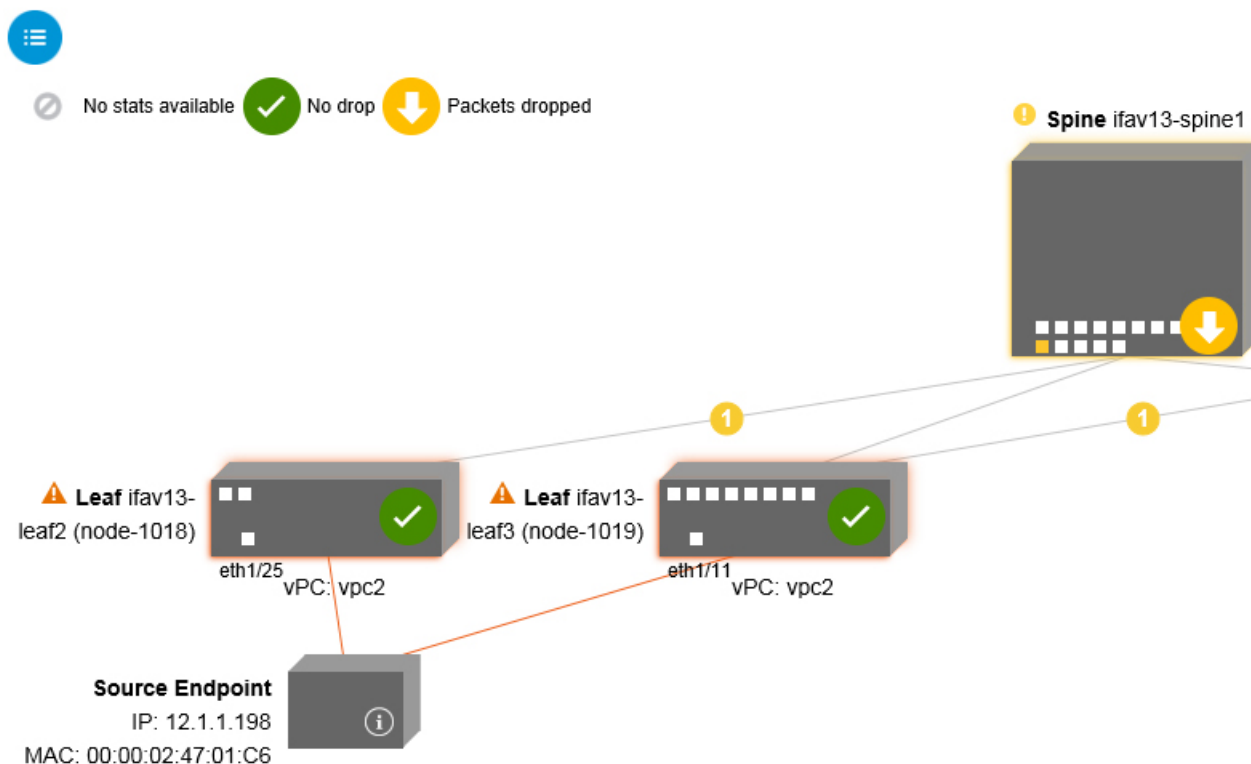
Related Topics

[Using the Drop/Statistics Troubleshooting Screen](#), on page 93

Using the Drop/Statistics Troubleshooting Screen

Click **Drop/Stats** in the **Navigation** pane to begin using the **Drop/Stats** troubleshooting screen.

The **Drop/Stats** window displays the topology with all the statistics from the drops so that you can clearly see where drops exist or not. You can click on any drop image to see more information for analysis.



Once you click a drop image, there are three tabs at the top of the **Drop/Stats** screen, and the statistics shown are localized to that particular leaf or switch.

The three statistics tabs are:

- **DROP STATS**

This tab shows the statistics for drop counters. The packets that are dropped at various levels are shown here.



Note By default, counters with zero values are hidden but the user can choose to see all the values.

- **CONTRACT DROPS**

This tab shows a list of the contract drops that have occurred, which are individual packet logs (ACL logs), and shows information for each packet such as the **Source Interface**, **Source IP address**, **Source Port**, **Destination IP address**, **Destination Port**, and **Protocol**.




Note Not every packet is displayed here.

• TRAFFIC STATS

This tab shows the statistics that indicate ongoing traffic. These are how many packets have been transferring.



Note By default, counters with zero values are hidden but the user can choose to see all the values.

You can also view all of the statistics for all managed objects at once by clicking the All icon () located in the top, left corner of the screen.

You also have the option to pick zero or non-zero drops. Checking the box for **Show stats with zero values** (located in the top, left corner of the screen) allows you to see all the existing drops. The fields for **Time**, **Affected Object**, **Stats**, and **Value** become populated with data for all the zero values.

If you do not check the **Show stats with zero values** box, you will see results with non-zero drops.



Note The same logic applies if you click the **All** icon. All three tabs (**DROP STATS**, **CONTRACT DROPS**, and **TRAFFIC STATS**) are also available and have the same type of information appearing.

Related Topics

[Using the Contracts Troubleshooting Screen](#), on page 94

Using the Contracts Troubleshooting Screen

Click **Contracts** in the **Navigation** pane to begin using the **Contracts** troubleshooting screen.

The **Contracts** troubleshooting screen displays the contracts that are applicable from the Source to the Destination and from the Destination to the Source.

Each one of the blue table heading rows indicates a filter. There are multiple rows under each filter that indicate multiple filter entries (**Protocol**, **L4 Src**, **L4 Dest**, **TCP Flags**, **Action**, **Nodes**, and **Hits**) for a particular leaf or switch.


Hovering over the certificate icon, shows you the contract name and the contract filter name. The text appearing on the right side of each blue table heading row (or filter) tells what type of contract it is, for example:

- Epg to Epg
- BD Allow
- Any to Any
- Context Deny

These contracts are categorized from the Source to the Destination and from the Destination to the Source.



Note The hits shown for each filter are cumulative (that is, the total hits for that contract hit, contract filter, or rule are shown for each particular leaf.) Statistics are refreshed automatically every (one) minute.

You can get policy information by hovering over the Information () icon. You can also see which EPGs are being referred to.



Note If there are no contracts between the endpoints, this will be indicated with a **There is no contract data** pop-up.

Related Topics

[Using the Events Troubleshooting Screen](#), on page 95


Using the Events Troubleshooting Screen

Click **Events and Audits** in the **Navigation** pane to begin using the **Events and Audits** troubleshooting screen.

If you click on an individual leaf or spine switch, you can see more detailed information about that individual event.

There are two tabs available: **EVENTS** and **DEPLOYMENT RECORDS**.

- **EVENTS** show event records for any changes that have occurred in systems (such as physical interfaces or VLANs, for example). There are individual events listed for each particular leaf. You can sort these events based on **Severity**, **Affected Object**, **Creation Time**, **Cause**, and **Description**.
- **DEPLOYMENT RECORDS** show the deployment of policies on physical interfaces, VLANs, VXLANs, and L3 CTXs. These records show the time when a VLAN was placed on a leaf because of the epg.

If you click the **All** icon () for the **All Changes** screen, you can see all the events indicating any changes that have occurred during your specified time interval (or troubleshooting session).

There are three tabs in the **All Changes** screen, including:

- **AUDITS**
Audits do not have a leaf association, which is why they are only available in the **All Changes** screen.
- **EVENTS** (described above)
- **DEPLOYMENT RECORDS** (described above)

Related Topics

[Using the Traceroute Troubleshooting Screen](#), on page 95

Using the Traceroute Troubleshooting Screen

Click **Traceroute** in the **Navigation** pane to begin using the **Traceroute** troubleshooting screen.

To create and run a traceroute for troubleshooting:

1. In the **TRACEROUTE** dialog box, choose a destination port from the **Destination Port** drop-down menu.
2. Choose a protocol from the **Protocol** pull-down menu. The options supported include:
 - **icmp**—This protocol is uni-directional, in that it does a traceroute from the Source leaf to the Destination endpoint only.
 - **tcp**—This protocol is also bi-directional, as described above for the **udp** protocol.
 - **udp**—This protocol is bi-directional, in that it does a traceroute from the Source leaf to the Destination endpoint, then from the Destination leaf back to the Source endpoint.



Note UDP, TCP and ICMP are the only supported protocols for IPv4. For IPv6, only UDP is supported.

3. Once you create a traceroute, click the **Play** (or Start) button to start the traceroute.



Note When you press the **Play** button, the policies are created on the system and a **Warning** message appears.

4. Click **OK** to proceed and the traceroute starts to run.
5. Click the **Stop** button to end the traceroute.



Note When you press the **Stop** button, the policies are removed from the system.

Once the traceroute completes, you can see where it was launched and what the result was. There is a pull-down menu next to **Traceroute Results** that shows where the traceroute was launched (from the Source to the Destination or from the Destination to the Source).

The result is also shown in the **Traceroute** dialog, which includes information for **Running Time**, **Traceroute Status**, **Destination Port**, and **Protocol**.

The results are represented by green and/or red arrows. A green arrow is used to represent each node in the path that responded to the traceroute probes. The beginning of a red arrow represents where the path ends as that's the last node that responded to the traceroute probes. You don't choose which direction to launch the traceroute. Instead, the traceroute is always started for the session. If the session is:

- EP to external IP or external IP to EP, the traceroute is always launched from EP to external IP.
- EP to EP and protocol is ICMP, the traceroute is always launched from the source to the destination.
- EP to EP and protocol is UDP/TCP, the traceroute is always bidirectional.



-
- Note**
- The **Traceroute Results** drop-down menu can be used to expose/visualize the results for each direction for scenario #3 above. In scenarios #1 and #2, it's always greyed out.
 - If the **Traceroute Status** shows as incomplete, this means you are still waiting for part of the data to come back. If the **Traceroute Status** shows as **complete**, then it is actually complete.
-

Related Topics

[Using the Atomic Counter Troubleshooting Screen](#), on page 97

Using the Atomic Counter Troubleshooting Screen

Click **Atomic Counter** in the **Navigation** pane to begin using the **Atomic Counter** troubleshooting screen.

The Atomic Counter screen is used to take source and destination information and create a counter policy based on that. You can create an atomic counter policy between two endpoints and monitor the traffic going back and forth from the Source to the Destination and from the Destination to the Source. You can determine how much traffic is going through and especially determine if any anomalies (drops or excess packets) are reported between the source and destination leaves.

There are **Play** (or Start) and **Stop** buttons at the top of the screen so that you can start and stop the atomic counter policy at any point and can count the packets that are being sent.



-
- Note** When you press the **Play** button, the policies are created on the system and the packet counter starts. When you press the **Stop** button, the policies are removed from the system.
-

The results are shown in two different formats. You can view them in either a brief format, which includes a summary, or in a longer format (by clicking on the **Expand** button). Both brief and expanded formats show both directions. The expanded format shows the cumulative counts plus the counts for each of the latest 30s intervals, while the brief format only shows the counts for cumulative and last interval.

Related Topics

[Using the SPAN Troubleshooting Screen](#), on page 97

Using the SPAN Troubleshooting Screen

Click **SPAN** in the **Navigation** pane to begin using the **SPAN** troubleshooting screen.

Using this screen, you can span (or mirror) bi-directional traffic and redirect it to the analyzer. In a SPAN session, you are making a copy and sending it to the analyzer.

This copy goes to a particular host (the analyzer IP address) and then you can use a software tool such as Wireshark to view the packets. The session information has source and destination information, session type, and the timestamp range.



-
- Note** When you press the **Play** button, the policies are created on the system. When you press the **Stop** button, the policies are removed from the system.
-



Note For a list of Troubleshooting Wizard CLI commands, see the *Cisco APIC Command-Line Interface User Guide*.

Creating a SPAN Session Using the Cisco APIC Troubleshooting CLI

This section demonstrates how to use the Cisco APIC troubleshooting CLIs to create a SPAN session.

Procedure

Step 1 **troubleshoot node session** <session_name> **nodename** <node_id>

To create a node-level session (global drop):

Example:

```
apic1(config)# troubleshoot node session 301-GD-APIC nodeid 301
```

Step 2 **troubleshoot node session** <session_name> **nodename** <node_id> **interface ethernet** <interface>

To create an interface-level session:

Example:

```
apic1(config)# troubleshoot node session 301-GD-APIC nodeid 301 interface eth1/3
```

Step 3 **troubleshoot node session** <session_name> **monitor destination** *apic_ip* **srcipprefix** <ip_prefix> **drop enable erspan-id**[optional]

To specify the destination as Cisco APIC and enable SPAN on drop:

Example:

```
apic1(config)# troubleshoot node session 301-GD-APIC monitor destination apic srcipprefix 13.13.13.13 drop enable
```

Step 4 **troubleshoot node session** <session_name> **monitor destination tenant** *tenant* **application** <app> **destip** <dest_ip>**srcipprefix**<ip_prefix>**drop enable erspan-id**[optional]

To specify an ERSPAN destination and enable SPAN on drop:

Example:

```
apic1(config)# troubleshoot node session 301-GD-APIC monitor destination tenant ERSPAN application A1 epg E1 destip 179.10.10.179 srcipprefix 31.31.13.31 drop enable
```

To check the SPAN-on-drop packets on the Cisco APIC when it is set as destination:

a. Disable the SPAN-on-drop session:

```
apic1(config)# no troubleshoot node session 301-GD-APIC monitor
```

b. Go to the drop-stats directory and check the DropPackets_*.pcap file:

```
/data2/techsupport/troubleshoot/node/Session_name/span_capture/drop-stats/DropPackets_*.pcap
```

L4 - L7 Services Validated Scenarios

The Troubleshooting Wizard allows you to provide two endpoints and see the corresponding topology between those endpoints. When L4 - L7 services exist between the two endpoints in the topology, you are able to view these as well.

This section describes the L4 - L7 scenarios that have been validated for this release. Within the L4 - L7 services, the number of topologies is very high, which means that you can have different configurations for firewalls, load balancers, and combinations of each. If a firewall exists between the two endpoints in the topology, the Troubleshooting Wizard retrieves the firewall data and connectivity from the firewall to the leafs. If a load balancer exists between the two endpoints, you can retrieve and view information up to the load balancer but not up to the server.

The following table shows the L4 - L7 service scenarios that were validated for the Troubleshooting Wizard:

Scenario	1	2	3	4	5	6
Number of Nodes	1	1	2	1	1	2
Device	GoTo FW (vrf split)	GoTo SLB	GoTo,GoTo FW,SLB	FW-GoThrough	SLB-GoTo	FW, SLB (GoThrough, GoTo)
Number of Arms	2	2	2	2	2	2
Consumer	EPG	EPG	EPG	L3Out	L3Out	L3Out
Provider	EPG	EPG	EPG	EPG	EPG	EPG
Device Type	VM	VM	VM	physical	physical	physical
Contract Scope	tenant	context	context	context	context	global
Connector Mode	L2	L2	L2, L2	L3, L2	L3	L3 / L2,L3
Service Attach	BSW	BSW	DL/PC	regular port	vPC	regular port
Client Attach	FEX	FEX	FEX	Regular Port	Regular Port	regular port
Server Attach	vPC	vPC	vPC	regular port	regular port	regular port

List of APIs for Endpoint to Endpoint Connections

The following is a list of the available Troubleshooting Wizard APIs for EP to EP (endpoint to endpoint) connections:

- [interactive API](#), on page 100
- [createsession API](#), on page 101
- [modifysession API](#), on page 102
- [atomiccounter API](#), on page 102
- [traceroute API](#), on page 103
- [span API](#), on page 103

- [generatereport API](#), on page 104
- [schedulesreport API](#), on page 104
- [getreportstatus API](#), on page 105
- [getreportslist API](#), on page 105
- [getsessionslist API](#), on page 106
- [getsessiondetail API](#), on page 106
- [deletesession API](#), on page 106
- [clearreports API](#), on page 107
- [contracts API](#), on page 107

interactive API

To create an endpoint (ep) to endpoint interactive troubleshooting session, use the **interactive** API. The module name is **troubleshoot.eptoeutils.topo** and the function is **getTopo**. The required argument (**req_args**) for the interactive API is - **session**.

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax	Description	Optional Arguments (opt_args)	Description
		- srcep	Source endpoint name
		- dstep	Destination endpoint name
		- srcip	Source endpoint IP address
		- dstip	Destination endpoint IP address
		- srcmac	Source endpoint MAC
		- dstmac	Destination endpoint MAC
		- srcextip	L3 external source IP address
		- dstextip	L3 external destination IP address
		- starttime	Start time of the troubleshooting session
		- endtime	End time of the troubleshooting session
		- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
		- description	Description about the session
		- scheduler	Scheduler name for report generation
		- srcepid	Obsolete
		- dstepid	Obsolete

- include	Obsolete
- format	Format of report to be generated
- ui	Used internally (ignore)
- sessionurl	Location of the report
-action	Start/stop/status etc. for traceroute/atomiccounter
- mode	Used internally
- _dc	Used internally
- ctx	Used internally

createsession API

To create an endpoint (ep) to endpoint troubleshooting session, use the **createsession** API. The module name is **troubleshoot.eptoeputils.session** and the function is **createSession**.

The required argument (**req_args**) for the createsession API is **- session** (session name).

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax	Optional Arguments (opt_args)	Description
	- srcep	Source endpoint name
	- dstep	Destination endpoint name
	- srcip	Source endpoint IP address
	- dstip	Destination endpoint IP address
	- srcmac	Source endpoint MAC
	- dstmac	Destination endpoint MAC
	- srcextip	L3 external source IP address
	- dstextip	L3 external destination IP address
	- starttime	Start time of the troubleshooting session
	- endtime	End time of the troubleshooting session
	- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
	- description	Description about the session
	- format	Format of report to be generated
	- ui	Used internally (ignore)

-action	Start/stop/status etc. for traceroute/atomiccounter
- scheduler	
- srctenant	Name of the tenant for the source endpoint
- srcapp	Name of the app for the source endpoint
- srcepg	Name of the endpoint group for the source endpoint
- dsttenant	Name of the tenant for the destination endpoint
- dstapp	Name of the app for the destination endpoint
- dstepg	Name of the endpoint group for the destination endpoint
- mode	Used internally

modifysession API

To modify an endpoint (ep) to endpoint troubleshooting session, use the **modifysession** API. The module name is **troubleshoot.eptoeutils.topo** and the function is **modifySession**.

The required arguments (**req_args**) for the modifysession API are **- session** (session name) and **- mode**.

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- starttime	Start time of the troubleshooting session
	- endtime	End time of the troubleshooting session
	- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
	- description	Description about the session

atomiccounter API

To create an endpoint (ep) to endpoint atomic counter session, use the **atomiccounter** API. The module name is **troubleshoot.eptoeutils.atomiccounter** and the function is **manageAtomicCounterPols**.

The required arguments (**req_args**) for the atomiccounter API include:

- - session
- - action
- - mode



Note There are no optional arguments (**opt_args**) for the atomiccounter API.

traceroute API

To create an endpoint (ep) to endpoint traceroute session using the API, use the **traceroute** API. The module name is **troubleshoot.eptoeputils.traceroute** and the function is **manageTraceroutePols**.

The required arguments (**req_args**) for the traceroute API include:

- - session (session name)
- - action (start/stop/status)
- - mode

Syntax Description	Optional Arguments (opt_args)	Description
	- protocol	Protocol name
	- dstport	Destination port name

span API

To create an endpoint (ep) to endpoint span troubleshooting session, use the **span** API. The module name is **troubleshoot.eptoeputils.span** and the function is **monitor**.

The required arguments (**req_args**) for the span API include:

- - session (session name)
- - action (start/stop/status)

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- srcep	Source endpoint name
	- dstep	Destination endpoint name
	- srcip	Source endpoint IP address
	- dstip	Destination endpoint IP address
	- sremac	Source endpoint MAC
	- dstmac	Destination endpoint MAC
	- srceextip	L3 external source IP address
	- dstextip	L3 external destination IP address
	- starttime	Start time of the troubleshooting session
	- endtime	End time of the troubleshooting session
	- latestmin	Time window for the troubleshooting session starting from start time (in minutes)

- description	Description about the session
- scheduler	Scheduler name for report generation
- srcepid	Obsolete
- dstepid	Obsolete
- include	Obsolete
- format	Format of report to be generated
- ui	Used internally (ignore)
- sessionurl	Location of the report
-action	Start/stop/status etc. for traceroute/atomiccounter
- srctenant	Name of the tenant for the source endpoint
- srcapp	Name of the app for the source endpoint
- srcepg	Name of the endpoint group for the source endpoint
- dsttenant	Name of the tenant for the destination endpoint
- dstapp	Name of the app for the destination endpoint
- dstepg	Name of the endpoint group for the destination endpoint
- mode	Used internally
- _dc	Used internally
- ctx	Used internally

generatereport API

To generate a troubleshooting report using the API, use the **generatereport** API. The module name is **troubleshoot.eptoeutils.report** and the function is **generateReport**.

The required arguments (**req_args**) for the generatereport API are **- session** (session name) and **- mode**.

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax	Description
Optional Arguments (opt_args)	Description
- include	Obsolete
- format	Format of report to be generated

schedulereport API

To schedule the generation of a troubleshooting report using the API, use the **schedulereport** API. The module name is **troubleshoot.eptoeutils.report** and the function is **scheduleReport**. The required argument (**req_args**) for the schedulereport API is **- session**

The required arguments (**req_args**) for the schedulereport API include:

- - session (session name)
- - scheduler (scheduler name)
- - mode

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- starttime	Start time of the troubleshooting session
	- endtime	End time of the troubleshooting session
	- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
	- include	Obsolete
	- format	Format of report to be generated
	- action	Start/stop/status etc. for traceroute/atomiccounter

getreportstatus API

To get the status of a generated report using the API, use the **getreportstatus** API. The module name is **troubleshoot.eptoeputils.report** and the function is **getStatus**.

The required arguments (**req_args**) for the getreportstatus API include:

- - session (session name)
- - sessionurl (session URL)
- - mode



Note There are no optional arguments (**opt_args**) for the getreportstatus API.

getreportslist API

To get a list of generated reports using the API, use the **getreportslist** API. The module name is **troubleshoot.eptoeputils.report** and the function is **getReportsList**.

The required arguments (**req_args**) for the getreportslist API are **- session** (session name) and **- mode**.



Note There are no optional arguments (**opt_args**) for the getreportslist API.

getsessionslist API

To get a list of troubleshooting sessions using the API, use the **getsessionslist** API. The module name is **troubleshoot.eptoeputils.session** and the function is **getSessions**.

The required argument (**req_args**) for the getsessionlist API is - **mode**.



Note There are no optional arguments (**opt_args**) for the getsessionlist API.

getsessiondetail API

To get specific details about a troubleshooting session using the API, use the **getsessiondetail** API. The module name is **troubleshoot.eptoeputils.session** and the function is **getSessionDetail**.

The required arguments (**req_args**) for the getsessiondetail API are - **session** (session name) and - **mode**.



Note There are no optional arguments (**opt_args**) for the getsessiondetail API.

deletesession API

To delete a particular troubleshooting session using the API, use the **deletesession** API. The module name is **troubleshoot.eptoeputils.session** and the function is **deleteSession**.

The required argument (**req_args**) for the deletesession API is - **session** (session name).

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax	Description
Optional Arguments (opt_args)	Description
- srcep	Source endpoint name
- dstep	Destination endpoint name
- srcip	Source endpoint IP address
- dstip	Destination endpoint IP address
- srcmac	Source endpoint MAC
- dstmac	Destination endpoint MAC
- srcextip	L3 external source IP address
- dstextip	L3 external destination IP address
- starttime	Start time of the troubleshooting session
- endtime	End time of the troubleshooting session
- latestmin	Time window for the troubleshooting session starting from start time (in minutes)

- description	Description about the session
- scheduler	Scheduler name for report generation
- srcepid	Obsolete
- dstepid	Obsolete
- include	Obsolete
- format	Format of report to be generated
- ui	Used internally (ignore)
- sessionurl	Location of report
- action	Start/stop/status etc. for traceroute/atomiccounter
- mode	Used internally
- _dc	Used internally
- ctx	Used internally

clearreports API

To clear the list of generated reports using the API, use the **clearreports** API. The module name is **troubleshoot.eptoeputils.report** and the function is **clearReports**.

The required arguments (**req_args**) for the clearreports API are- **session** (session name) and - **mode**.



Note There are no optional arguments (**opt_args**) for the clearreports API.

contracts API

To get contracts information using the API, use the **contracts** API. The module name is **troubleshoot.eptoeputils.contracts** and the function is **getContracts**.

The required arguments (**req_args**) for the contracts API are- **session** (session name) and -**mode**.

There are no optional arguments (**opt_args**) for the contracts API.

List of APIs for Endpoint to Layer 3 External Connections

The following is a list of the available Troubleshooting Wizard APIs for EP to EP (endpoint to endpoint) connections:

- [interactive API, on page 108](#)
- [modifysession API, on page 109](#)
- [atomiccounter API, on page 110](#)
- [traceroute API, on page 111](#)

- [span API](#), on page 112
- [generatereport API](#), on page 112
- [schedulesreport API](#), on page 113
- [getreportstatus API](#), on page 105
- [getreportslist API](#), on page 105
- [clearreports API](#), on page 107
- [createsession API](#), on page 108
- [getsessionslist API](#), on page 115
- [getsessiondetail API](#), on page 116
- [deletesession API](#), on page 117
- [contracts API](#), on page 117
- [ratelimit API](#), on page 118
- [13ext API](#), on page 119

interactive API

To create an endpoint (ep) to Layer 3 (L3) external interactive troubleshooting session, use the **interactive** API. The module name is **troubleshoot.epextutils.epext_topo** and the function is **getTopo**. The required arguments (**req_args**) for the interactive API are **- session**, **- include**, and **- mode**.

The following table shows the optional argument (**opt_args**):

Syntax Description	Optional Arguments (opt_args)	Description
	- refresh	

createsession API

To create an endpoint (Ep) to Layer 3 (L3) external troubleshooting session using the API, use the **createsession** API. The module name is **troubleshoot.epextutils.epextsession** and the function is **createSession**. The required argument (**req_args**) for the createsession API is **- session** (session name).

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- srecep	Source endpoint name
	- dstep	Destination endpoint name
	- srcip	Source endpoint IP address
	- dstip	Destination endpoint IP address
	- srcmac	Source endpoint MAC

- dstmac	Destination endpoint MAC
- srcextip	L3 external source IP address
- dstextip	L3 external destination IP address
- starttime	Start time of the troubleshooting session
- endtime	End time of the troubleshooting session
- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
- description	Description about the session
- scheduler	Scheduler name for report generation
- srcepid	Obsolete
- dstepid	Obsolete
- include	Obsolete
- format	Format of report to be generated
- ui	Used internally (ignore)
- sessionurl	Location of the report
-action	Start/stop/status etc. for traceroute/atomiccounter
- mode	Used internally
- _dc	Used internally
- ctx	Used internally

modifysession API

To modify an endpoint (Ep) to Layer 3 (L3) external troubleshooting session, use the **modifysession** API. The module name is **troubleshoot.epextutils.epextsession** and the function is **modifySession**. The required argument (**req_args**) for the modifysession API is **- session** (session name).

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- srcep	Source endpoint name
	- dstep	Destination endpoint name
	- srcip	Source endpoint IP address
	- dstip	Destination endpoint IP address
	- srcmac	Source endpoint MAC

- dstmac	Destination endpoint MAC
- srcextip	L3 external source IP address
- dstextip	L3 external destination IP address
- starttime	Start time of the troubleshooting session
- endtime	End time of the troubleshooting session
- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
- description	Description about the session
- scheduler	Scheduler name for report generation
- srcepid	Obsolete
- dstepid	Obsolete
- include	Obsolete
- format	Format of report to be generated
- ui	Used internally (ignore)
- sessionurl	Location of the report
-action	Start/stop/status etc. for traceroute/atomiccounter
- mode	Used internally
- _dc	Used internally
- ctx	Used internally

atomiccounter API

To create an endpoint (ep) to endpoint atomic counter session, use the **atomiccounter** API. The module name is **troubleshoot.epextutils.epext_ac** and the function is **manageAtomicCounterPols**.

The required arguments (**req_args**) for the atomiccounter API include:

- - session (session name)
- - action (start/stop/status)

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- srcep	Source endpoint name
	- dstep	Destination endpoint name

- srcip	Source endpoint IP address
- dstip	Destination endpoint IP address
- srcmac	Source endpoint MAC
- dstmac	Destination endpoint MAC
- srcextip	L3 external source IP address
- dstextip	L3 external destination IP address
- starttime	Start time of the troubleshooting session
- endtime	End time of the troubleshooting session
- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
- ui	Used internally (ignore)
- mode	Used internally
- _dc	Used internally
- ctx	Used internally

traceroute API

To create an endpoint (ep) to Layer 3 external traceroute troubleshooting session using the API, use the **traceroute** API. The module name is **troubleshoot.epextutils.epext_traceroute** and the function is **manageTraceroutePols**.

The required arguments (**req_args**) for the traceroute API include:

- - session (session name)
- - action (start/stop/status)

Syntax Description

Optional Arguments (opt_args)	Description
- protocol	Protocol name
- dstport	Destination port name
- srcep	Source endpoint
- dstep	Destination endpoint
- srcip	Source IP address
- dstip	Destination IP address
- srcextip	Source external IP address
- dstlp	Destination external IP address

- ui	Used internally (ignore)
- mode	Used internally
- _dc	Used internally
- ctx	Used internally

span API

To create an endpoint (Ep) to Layer 3 (L3) external span troubleshooting session, use the **span** API. The module name is **troubleshoot.epextutils.epext_span** and the function is **monitor**.

The required arguments (**req_args**) for the span API include:

- - session (session name)
- - action (start/stop/status)
- - mode

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- portslst	List of ports
	- dstapic	Destination APIC
	- srcipprefix	Source endpoint IP address prefix
	- flowid	Flow ID
	- dstepg	Destination endpoint group
	- dstip	Destination endpoint IP address
	- analyser	???
	- desttype	Destination type
	- spansrports	Span source ports

generatereport API

To generate a troubleshooting report using the API, use the **generatereport** API. The module name is **troubleshoot.eptoeutils.report** and the function is **generateReport**.

The required argument (**req_args**) for the generatereport API is - **session** (session name).

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- srcep	Source endpoint name

- dstep	Destination endpoint name
- srcip	Source endpoint IP address
- dstip	Destination endpoint IP address
- srcmac	Source endpoint MAC
- dstmac	Destination endpoint MAC
- srcextip	L3 external source IP address
- dstextip	L3 external destination IP address
- starttime	Start time of the troubleshooting session
- endtime	End time of the troubleshooting session
- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
- description	Description about the session
- scheduler	Scheduler name for report generation
- srcepid	Obsolete
- dstepid	Obsolete
- include	Obsolete
- format	Format of report to be generated
- ui	Used internally (ignore)
- sessionurl	Location of the report
-action	Start/stop/status etc. for traceroute/atomiccounter
- mode	Used internally
- _dc	Used internally
- ctx	Used internally

schedulingreport API

To schedule the generation of a troubleshooting report using the API, use the **schedulingreport** API. The module name is **troubleshoot.eptoeutils.report** and the function is **scheduleReport**. The required argument (**req_args**) for the schedulingreport API is - **session**

The required arguments (**req_args**) for the schedulingreport API include:

- - session (session name)
- - scheduler (scheduler name)

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- srcep	Source endpoint
	- dstep	Destination endpoint
	- srcip	Source endpoint IP address
	- dstip	Destination endpoint IP address
	- srcmac	Source endpoint MAC
	- dstmac	Destination endpoint MAC
	- srcextip	L3 external source IP address
	- dstextip	L3 external destination IP address
	- starttime	Start time of the troubleshooting session
	- endtime	End time of the troubleshooting session
	- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
	- description	Description about the session
	- srecepid	Obsolete
	- dstepid	Obsolete
	- include	Obsolete
	- format	Format of report to be generated
	- ui	Used internally (ignore)
	- sessionurl	Location of the report
	-action	Start/stop/status etc. for traceroute/atomiccounter
	- mode	Used internally
	- _dc	Used internally
	- ctx	Used internally

getreportstatus API

To get the status of a generated report using the API, use the **getreportstatus** API. The module name is **troubleshoot.eptoeputils.report** and the function is **getStatus**.

The required arguments (**req_args**) for the getreportstatus API include:

- - session (session name)

- - sessionurl (session URL)
- - mode



Note There are no optional arguments (**opt_args**) for the getreportstatus API.

getreportslist API

To get a list of generated reports using the API, use the **getreportslist** API. The module name is **troubleshoot.epptoeputils.report** and the function is **getReportsList**.

The required arguments (**req_args**) for the getreportslist API are- **session** (session name) and - **mode**.



Note There are no optional arguments (**opt_args**) for the getreportslist API.

getsessionslist API

To get a list of troubleshooting sessions using the API, use the **getsessionslist** API. The module name is **troubleshoot.epextutils.epextsession** and the function is **getSessions**.



Note There are no required arguments for this API.

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax	Description	Optional Arguments (opt_args)	Description
		- session	Session name
		- srcep	Source endpoint name
		- dstep	Destination endpoint name
		- srcip	Source endpoint IP address
		- dstip	Destination endpoint IP address
		- srcmac	Source endpoint MAC
		- dstmac	Destination endpoint MAC
		- srcextip	L3 external source IP address
		- dstextip	L3 external destination IP address
		- starttime	Start time of the troubleshooting session
		- endtime	End time of the troubleshooting session

- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
- description	Description about the session
- scheduler	Scheduler name for report generation
- srcepid	Obsolete
- dstepid	Obsolete
- include	Obsolete
- format	Format of report to be generated
- ui	Used internally (ignore)
- sessionurl	Location of report
- action	Start/stop/status etc. for traceroute/atomiccounter
- mode	Used internally
- _dc	Used internally
- ctx	Used internally

getsessiondetail API

To get specific details about a troubleshooting session using the API, use the **getsessiondetail** API. The module name is **troubleshoot.epextutils.session** and the function is **getSessionDetail**. The required argument (**req_args**) for the **getsessiondetail** API is **- session** (session name).

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- srcep	Source endpoint name
	- dstep	Destination endpoint name
	- srcip	Source endpoint IP address
	- dstip	Destination endpoint IP address
	- srcmac	Source endpoint MAC
	- dstmac	Destination endpoint MAC
	- srcextip	L3 external source IP address
	- dstextip	L3 external destination IP address
	- starttime	Start time of the troubleshooting session
	- endtime	End time of the troubleshooting session

- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
- description	Description about the session
- scheduler	Scheduler name for report generation
- srcepid	Obsolete
- dstepid	Obsolete
- include	Obsolete
- format	Format of report to be generated
- ui	Used internally (ignore)
- sessionurl	Location of report
- action	Start/stop/status etc. for traceroute/atomiccounter
- mode	Used internally
- _dc	Used internally
- ctx	Used internally

deletesession API

To delete a particular troubleshooting session using the API, use the **deletesession** API. The module name is **troubleshoot.epextutils.epextsession** and the function is **deleteSession**.

The required arguments (**req_args**) for the deletesession API are **- session** (session name) and **- mode**.



Note There are no optional arguments (**opt_args**) for the deletesession API.

clearreports API

To clear the list of generated reports using the API, use the **clearreports** API. The module name is **troubleshoot.eptoeputils.report** and the function is **clearReports**.

The required arguments (**req_args**) for the clearreports API are **- session** (session name) and **- mode**.



Note There are no optional arguments (**opt_args**) for the clearreports API.

contracts API

To get contracts information using the API, use the **contracts** API. The module name is **troubleshoot.epextutils.epext_contracts** and the function is **getContracts**. The required argument (**req_args**) for the contracts API is **- session** (session name).

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- srcep	Source endpoint name
	- dstep	Destination endpoint name
	- srcip	Source endpoint IP address
	- dstip	Destination endpoint IP address
	- srcmac	Source endpoint MAC
	- dstmac	Destination endpoint MAC
	- srcextip	L3 external source IP address
	- dstextip	L3 external destination IP address
	- starttime	Start time of the troubleshooting session
	- endtime	End time of the troubleshooting session
	- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
	- epext	Endpoint to external
	- mode	Used internally
	- _dc	Used internally
	- ctx	Used internally
	- ui	Used internally (ignore)

ratelimit API

This section provides information on the the **ratelimit** API. The module name is **troubleshoot.eptoeputils.ratelimit** and the function is **control**. The required argument (**req_args**) for the ratelimit API is **- action** (start/stop/status).

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- srcep	Source endpoint name
	- dstep	Destination endpoint name
	- srcip	Source endpoint IP address
	- dstip	Destination endpoint IP address
	- srcmac	Source endpoint MAC

- dstmac	Destination endpoint MAC
- srcextip	L3 external source IP address
- dstextip	L3 external destination IP address
- starttime	Start time of the troubleshooting session
- endtime	End time of the troubleshooting session
- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
- epext	Endpoint to external
- mode	Used internally
- _dc	Used internally
- ctx	Used internally

13ext API

This section provides information on the the **13ext** API. The module name is **troubleshoot.epextutils.13ext** and the function is **execute**. The required argument (**req_args**) for the 13ext API is - **action** (start/stop/status).

The following table lists the optional arguments (**opt_args**) and descriptions for each.

Syntax Description	Optional Arguments (opt_args)	Description
	- srcep	Source endpoint name
	- dstep	Destination endpoint name
	- srcip	Source endpoint IP address
	- dstip	Destination endpoint IP address
	- srcmac	Source endpoint MAC
	- dstmac	Destination endpoint MAC
	- srcextip	L3 external source IP address
	- dstextip	L3 external destination IP address
	- starttime	Start time of the troubleshooting session
	- endtime	End time of the troubleshooting session
	- latestmin	Time window for the troubleshooting session starting from start time (in minutes)
	- epext	Endpoint to external
	- mode	Used internally

Checking for Configuration Synchronization Issues

When you make a request in Cisco Application Centric Infrastructure (APIC)—for example, changing a configuration—you generally see immediately that the change has occurred. However, if you encounter an issue with Cisco APIC, you can check in the GUI to see if there are any transactions involving user-configurable objects that have yet to take effect. You can use information in the panel to help with debugging.

The **Configuration Objects Pending Resolution** panel in the Cisco APIC GUI tells you if there are delays.

Before you begin

Procedure

- Step 1** Log in to Cisco APIC.
 - Step 2** Click the settings icon (the gear symbol) in the upper right of the screen and choose **Config Sync Issues**.
 - Step 3** In the **Configuration Objects Pending Resolution** panel, check if anything is listed in the table.
If there are no entries in the table, there are no synchronization issues.
 - Step 4** If there are any entries, capture the information in the table and use it for debugging or working with Cisco support.
-

Viewing User Activities

In situations where an admin notices a change to the Cisco APIC setup, the admin can use the **User Activities** feature to view a 2-week history of actions performed by a user. The historical data includes a timestamp of when the action occurred, the user who performed the action, the action the user performed, the affected object, and a description.

Accessing User Activities

The **User Activities** window enables you to view a 2-week history of user activities performed in the Cisco APIC GUI.

Procedure

- Step 1** From the menu bar, choose **System > Active Sessions** .
The **Active Session** window appears.
 - Step 2** Right-click on an active session and choose **User Activities**.
A list of user activities appears.
- Note** For an explanation of a field, click the help icon in the top-right corner of the **Active Session** window to display the help file.

- Step 3** Click the **Actions in the last** drop-down menu to choose how far back in history you want to view the user activities.

About the Embedded Logic Analyzer Module

ELAM (Embedded Logic Analyzer Module) is an engineering tool that enables you to look inside Cisco ASICs and understand how a packet is being forwarded. ELAM is embedded within the forwarding pipeline and can capture a packet in real time without affecting performance or control plane resources. ELAM can perform the following functions:

- Determine if a packet reached the forwarding engine
- Specify the port and VLAN of the packet that was received
- View the packet (Layer 2 to Layer 4 data)
- Check if the packet was altered where it was sent

Generating an ELAM Report in the Simplified Output for Modular Switches

The Cisco Application Policy Infrastructure Controller (APIC) 4.2(1) release introduces simplified, human-readable ELAM output. Only switch models with EX, FX, or FX2 at the end of the switch name support the simplified output. Use the following procedure for modular switches.

Procedure

- Step 1** Run the ELAM tool to collect the packet forwarding information. The exact commands and parameters depend on your hardware.
- Step 2** Run the **ereport** command to create ELAM reports of the packet forwarding information in the original format and the simplified format.

Example:

```
module-1(DBG-elam-el6)# ereport
Python available. Continue ELAM decode with LC Pkg
ELAM REPORT

=====
                        Trigger/Basic Information
=====
ELAM Report File      : /tmp/logs/elam_2019-09-04-51m-13h-30s.txt
.
.
.

module-1(DBG-elam-el6)# exit
module-1(DBG-elam)# exit
module-1# exit

apic1-leaf11# cd /tmp/logs
apic1-leaf11# ls | grep elam
elam_2019-09-04-51m-13h-30s.txt
pretty_elam_2019-09-04-51m-13h-30s.txt
apic1-leaf11#
```

ELAM saves the output files in the `/tmp/logs/` directory. In the example, the `elam_2019-09-04-51m-13h-30s.txt` file is the ELAM report in the original format. The `pretty_elam_2019-09-04-51m-13h-30s.txt` file is the ELAM report in the simplified format. However, the simplified format file will be empty. You must perform additional steps to get the report in the simplified format.

Step 3 Upload the original format ELAM report to the `/bootflash` directory on the supervisor.
In the example, this report is the `elam_2019-09-04-51m-13h-30s.txt` file.

Step 4 Log in to the supervisor as admin.

Step 5 Change the directory to `/tmp`, or any directory with write privileges for the admin user.

Example:

```
# cd /tmp
```

Step 6 Run the `decode_elam_parser` command on the original format ELAM report.

Example:

```
# decode_elam_parser /bootflash/elam_2019-09-04-51m-13h-30s.txt
```

The `decode_elam_parser` command saves the simplified output file in the current directory.

Generating an ELAM Report in the Simplified Output for Fixed Form-Factor Switches

The Cisco Application Policy Infrastructure Controller (APIC) 4.2(1) release introduces simplified, human-readable ELAM output. Only switch models with EX, FX, or FX2 at the end of the switch name support the simplified output. Use the following procedure for fixed form-factor leaf switches and spine switches.

Procedure

Step 1 Run the ELAM tool to collect the packet forwarding information. The exact commands and parameters depend on your hardware.

Step 2 Run the `ereport` command to create ELAM reports of the packet forwarding information in the original format and the simplified format.

Example:

```
module-1(DBG-elam-insel6)# ereport
Python available. Continue ELAM decode with LC Pkg
ELAM REPORT

=====
                          Trigger/Basic Information
=====
ELAM Report File      : /tmp/logs/elam_2019-09-04-51m-13h-30s.txt
.
.
.

module-1(DBG-elam-insel6)# exit
module-1(DBG-elam)# exit
```

```

module-1# exit

apic1-leaf11# cd /tmp/logs
apic1-leaf11# ls | grep elam
elam_2019-09-04-51m-13h-30s.txt
pretty_elam_2019-09-04-51m-13h-30s.txt
apic1-leaf11#

```

ELAM saves the output files in the /tmp/logs/ directory. In the example, the elam_2019-09-04-51m-13h-30s.txt file is the ELAM report in the original format. The pretty_elam_2019-09-04-51m-13h-30s.txt file is the ELAM report in the simplified format.

acidiag Command

To troubleshoot operations on the Cisco APIC, use the **acidiag** command.



Caution This command is not intended for every day operation of ACI. Running all forms of the command can be very disruptive and cause major issues in your network if not used properly. Make sure you understand the full effect on your fabric before running them.

Cluster Commands

```

acidiag
acidiag avread
acidiag fnvread
acidiag fnvreadex

```

Syntax Description	Option	Function
	avread	Displays APICs within the cluster. The avread output includes: <ul style="list-style-type: none"> • Cluster of —Operational cluster size • out of targeted—The desired cluster size • active= —Indicates whether the APIC is reachable • health= —The overall APIC health summary. Displays services with degraded health scores. • chassisID= —The known chassis IDs for a given APIC. <p>Note Peer chassis IDs can be incorrect for APICs not currently in the cluster.</p>

Option	Function
bootcurr	On the next boot, the APIC system will boot the current APIC image in the Linux partition. This option is not expected to normally be used.
bootother	On the next boot, the APIC system will boot the previous APIC image in the Linux partition. This option is not expected to normally be used.
bond0test	Disruptive test of the APIC connection to the leaf. This is used for internal Cisco testing purposes only and outside of that could cause issues with the APIC connection to the fabric.
fnvread	Displays the address and state of switch nodes registered with the fabric.
fnvreadex	Displays additional information for switch nodes registered with the fabric.
linkflap	Brings down and back up a specified APIC interface.
preservelogs	APIC will archive current logs. During a normal reboot this automatically occurs. This option can be used prior to a hard reboot.
run	Two available options are iptables-list and lldptool. The iptables-list is used to display the Linux iptables, which are controlled by the mgmt Tenant contracts. lldptool is used to display lldp information which is sent or received by the APIC.
rvread	Summarizes the data layer state. The output shows a summary of the data layer state for each service. The shard view shows replicas in ascending order.
acidiag rvread <i>service</i>	Displays the data layer state for a service on all shards across all replicas. Note For an example, see Examples, on page 128
acidiag rvread <i>service shard</i>	Displays the data layer state for a service on a specific shard across all replicas. Note For an example, see Examples, on page 128
acidiag rvread <i>service shard replica</i>	Displays the data layer state for a service on a specific shard and replica. Note For an example, see Examples, on page 128

Option	Function
validateimage	Prior to loading an image into the firmware repository, the image can be validated. Note that this function runs as a normal part of the process of the image being added into the repository.
validateenginxconf	Validates the generated nginx configuration file on APIC to ensure nginx can start with that configuration file. This is meant for debug use, in cases where the nginx webserver is not running on APIC.

Service IDs

The service IDs listed in the table below are also visible when entering the **man acidiag** command.

Table 2: Service IDs

Service	ID
cliD	1
controller	2
eventmgr	3
extXMLApi	4
polycyelem	5
polycymgr	6
reader	7
ae	8
topomgr	9
observer	10
dbgr	11
observerelem	12
dbgrelem	13
vmmmgr	14
nxosmock	15
bootmgr	16
appliancedirector	17
adrelay	18

Service	ID
ospaagent	19
vleafelem	20
dhcpd	21
scripthandler	22
idmgr	23
ospaelem	24
osh	25
opflexagent	26
opflexelem	27
confelem	28
vtap	29
snmpd	30
opflexp	31
analytics	32
policydist	33
plghandler	34
domainmgr	35
licensemgr	36
N/A	37
platformmgr	38
edmgr	39

Table 3: Data States

State	ID
COMATOSE	0
NEWLY_BORN	1
UNKNOWN	2
DATA_LAYER_DIVERGED	11
DATA_LAYER_DEGRADED_LEADERSHIP	12

State	ID
DATA_LAYER_ENTIRELY_DIVERGED	111
DATA_LAYER_PARTIALLY_DIVERGED	112
DATA_LAYER_ENTIRELY_DEGRADED_LEADERSHIP	121
DATA_LAYER_PARTIALLY_DEGRADED_LEADERSHIP	122
FULLY_FIT	255

System Keywords

```
acidiag [start|stop|restart] [mgmt|xinetd]
```

```
acidiag installer -u imageurl -c
```

```
acidiag reboot
```

```
acidiag touch [clean|setup]
```

```
acidiag verifyapic
```

Syntax Description

Option	Function
-c	Specifies a clean install
-u	Specifies a URL for the APIC image.
<i>imageurl</i>	Specifies an APIC image.
installer	Installs a new image on the APIC, -c for clean install
mgmt	Specifies all services on the APIC.
reboot	Reboots the APIC.
restart	Restarts services on an APIC.
start	Starts services on an APIC.
stop	Stops services on an APIC.
touch [clean setup]	Resets the APIC configuration. <ul style="list-style-type: none"> • The clean option removes all policy data while retaining the APIC network configuration (such as fabric name, IP address, login) • The setup option removes both policy data and the APIC network configuration.
verifyapic	Displays the APIC software version.

Option	Function
xinetd	Specifies xinetd (extended internet daemon) service, which controls the ssh and telnet daemons. Beginning with the 6.0(2) release, telnet is not supported.

Diagnostic Keywords

acidiag crashsuspecttracker

acidiag dbgtoken

acidiag version

Syntax Description

Option	Function
crashsuspecttracker	Tracks states of a service or data subset that indicate a crash.
dbgtoken	Generates a token used to generate a root password. This is to be used as directed while working with the TAC as needed.
version	Displays the APIC ISO software version.

Examples

The following examples show how to use the **acidiag** command:

```
apic1# acidiag version 2.2.1o
```

```
apic1# acidiag verifyapic
openssl_check: certificate details
subject= CN=ABC12345678,serialNumber=PID:APIC-SERVER-L1 SN:ABC12345678
issuer= CN=Cisco Manufacturing CA,O=Cisco Systems
notBefore=Sep 28 17:17:42 2016 GMT
notAfter=Sep 28 17:27:42 2026 GMT
openssl_check: passed
ssh_check: passed
all_checks: passed
```

```
apic1# acidiag avread
Local appliance ID=1 ADDRESS=10.0.0.1 TEP ADDRESS=10.0.0.0/16 ROUTABLE IP ADDRESS=0.0.0.0
CHASSIS_ID=1009f750-adab-11e9-a044-8dbd212cd556
Cluster of 7 lm(t):1(2019-08-08T01:02:17.961-07:00) appliances (out of targeted 7
lm(t):7(2019-08-08T03:50:57.240-07:00)) with FABRIC_DOMAIN name=ACI Fabric1 set to
version=apic-4.2(0.235j) lm(t):1(2019-08-17T01:09:16.413-07:00); discoveryMode=PERMISSIVE
lm(t):0(1969-12-31T17:00:00.007-07:00); drrMode=OFF lm(t):0(1969-12-31T17:00:00.007-07:00);
kafkaMode=OFF lm(t):0(1969-12-31T17:00:00.007-07:00)
  appliance id=1 address=10.0.0.1 lm(t):1(2019-08-08T01:02:08.544-07:00) tep
address=10.0.0.0/16 lm(t):1(2019-08-08T01:02:08.544-07:00) routable address=0.0.0.0
lm(t):1(zeroTime) oob address=172.23.96.10/21 lm(t):1(2019-08-08T01:02:18.218-07:00)
version=4.2(0.235j) lm(t):1(2019-08-15T15:22:00.158-07:00)
chassisId=1009f750-adab-11e9-a044-8dbd212cd556 lm(t):1(2019-08-15T15:22:00.158-07:00)
capabilities=0X3EEFFFFFFFFF--0X2020--0X7F lm(t):1(2019-08-17T01:13:46.997-07:00)
```



```

rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-08T01:02:18.228-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-08T01:02:18.228-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-08T01:02:18.228-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-08T01:02:18.228-07:00)
cntrlSbst=(APPROVED, FCH1748V0SZ) lm(t):1(2019-08-15T15:22:00.158-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=1
lm(t):1(2019-08-08T01:02:08.544-07:00) commissioned=YES lm(t):1(zeroTime) registered=YES
lm(t):1(2019-08-08T01:02:08.544-07:00) standby=NO lm(t):1(2019-08-08T01:02:08.544-07:00)
DRR=NO lm(t):0(zeroTime) apicX=NO lm(t):1(2019-08-08T01:02:08.544-07:00) virtual=NO
lm(t):1(2019-08-08T01:02:08.544-07:00) active=YES(2019-08-08T01:02:08.544-07:00)
health=(applnc:255 lm(t):1(2019-08-17T01:39:26.296-07:00) svc's)
    appliance id=2 address=10.0.0.2 lm(t):7(2019-08-08T03:50:55.470-07:00) tep
address=10.0.0.0/16 lm(t):2(2019-07-23T17:51:38.997-07:00) routable address=0.0.0.0
lm(t):0(zeroTime) oob address=172.23.96.11/21 lm(t):1(2019-08-18T23:14:28.720-07:00)
version=4.2(0.235j) lm(t):2(2019-08-15T15:22:00.300-07:00)
chassisId=694e6a98-adac-11e9-ad79-d1f60e3ee822 lm(t):2(2019-08-15T15:22:00.300-07:00)
capabilities=0X3EEEEEEEEEE--0X2020--0X2 lm(t):2(2019-08-14T07:55:10.074-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.829-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.829-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.829-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.829-07:00)
cntrlSbst=(APPROVED, FCH1748V0MS) lm(t):2(2019-08-15T15:22:00.300-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=1
lm(t):2(2019-08-08T01:42:03.670-07:00) commissioned=YES lm(t):1(2019-08-08T01:02:17.961-07:00)
    registered=YES lm(t):7(2019-07-24T15:24:25.693-07:00) standby=NO
lm(t):2(2019-08-08T01:42:03.670-07:00) DRR=NO lm(t):1(2019-08-08T01:02:17.961-07:00) apicX=NO
    lm(t):2(2019-08-08T01:42:03.670-07:00) virtual=NO lm(t):0(zeroTime)
active=YES(2019-08-13T17:02:32.983-07:00) health=(applnc:255
lm(t):2(2019-08-17T01:32:51.454-07:00) svc's)
    appliance id=3 address=10.0.0.3 lm(t):7(2019-08-08T03:50:55.470-07:00) tep
address=10.0.0.0/16 lm(t):3(2019-07-23T19:05:56.405-07:00) routable address=0.0.0.0
lm(t):0(zeroTime) oob address=172.23.96.12/21 lm(t):1(2019-08-18T23:14:28.721-07:00)
version=4.2(0.235j) lm(t):3(2019-08-15T15:21:59.893-07:00)
chassisId=1f98b916-adb7-11e9-a6f8-abe00a04e8e6 lm(t):3(2019-08-15T15:21:59.893-07:00)
capabilities=0X3EEEEEEEEEE--0X2020--0X4 lm(t):3(2019-08-14T07:55:22.256-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
cntrlSbst=(APPROVED, FCH1930V1X6) lm(t):3(2019-08-15T15:21:59.893-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=1
lm(t):3(2019-08-08T02:15:20.560-07:00) commissioned=YES lm(t):2(2019-08-08T01:42:15.337-07:00)
    registered=YES lm(t):7(2019-07-24T15:24:25.693-07:00) standby=NO
lm(t):3(2019-08-08T02:15:20.560-07:00) DRR=NO lm(t):2(2019-08-08T01:42:15.337-07:00) apicX=NO
    lm(t):3(2019-08-08T02:15:20.560-07:00) virtual=NO lm(t):0(zeroTime)
active=YES(2019-08-13T17:02:33.182-07:00) health=(applnc:255
lm(t):3(2019-08-15T16:08:46.119-07:00) svc's)
    appliance id=4 address=10.0.0.4 lm(t):7(2019-08-08T03:50:55.470-07:00) tep
address=10.0.0.0/16 lm(t):4(2019-07-23T17:46:15.545-07:00) routable address=0.0.0.0
lm(t):0(zeroTime) oob address=172.23.97.231/21 lm(t):1(2019-08-18T23:14:28.717-07:00)
version=4.2(0.235j) lm(t):4(2019-08-15T15:22:00.669-07:00)
chassisId=3a7f38aa-adac-11e9-8869-a9e520cdc042 lm(t):4(2019-08-15T15:22:00.669-07:00)
capabilities=0X3EEEEEEEEEE--0X2020--0X8 lm(t):4(2019-08-14T07:54:59.490-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.825-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.825-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.825-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.825-07:00)
cntrlSbst=(APPROVED, FCH1902V1WW) lm(t):4(2019-08-15T15:22:00.669-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=1
lm(t):4(2019-08-08T02:40:09.610-07:00) commissioned=YES lm(t):3(2019-08-08T02:15:32.613-07:00)
    registered=YES lm(t):7(2019-07-24T15:24:25.693-07:00) standby=NO
lm(t):4(2019-08-08T02:40:09.610-07:00) DRR=NO lm(t):3(2019-08-08T02:15:32.613-07:00) apicX=NO
    lm(t):4(2019-08-08T02:40:09.610-07:00) virtual=NO lm(t):0(zeroTime)
active=YES(2019-08-15T15:21:59.914-07:00) health=(applnc:255

```

```

lm(t):4(2019-08-17T01:39:26.477-07:00) svc's)
  appliance id=5 address=10.0.0.5 lm(t):7(2019-08-08T03:50:55.470-07:00) tep
address=10.0.0.0/16 lm(t):5(2019-07-23T19:05:11.089-07:00) routable address=0.0.0.0
lm(t):0(zeroTime) oob address=172.23.97.232/21 lm(t):1(2019-08-18T23:14:28.723-07:00)
version=4.2(0.235j) lm(t):5(2019-08-15T15:22:00.248-07:00)
chassisId=35428666-adb7-11e9-a315-1d7671b518b3 lm(t):5(2019-08-15T15:22:00.248-07:00)
capabilities=0X3EFFFFFFFF--0X2020--0X10 lm(t):5(2019-08-14T07:55:19.573-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.854-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.854-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.854-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.854-07:00)
cntrlSbst=(APPROVED, FCH1902VLEG) lm(t):5(2019-08-15T15:22:00.248-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=1
lm(t):5(2019-08-08T03:03:50.338-07:00) commissioned=YES lm(t):4(2019-08-08T02:40:15.939-07:00)
  registered=YES lm(t):7(2019-07-24T15:24:25.693-07:00) standby=NO
lm(t):5(2019-08-08T03:03:50.338-07:00) DRR=NO lm(t):4(2019-08-08T02:40:15.939-07:00) apicX=NO
  lm(t):5(2019-08-08T03:03:50.338-07:00) virtual=NO lm(t):0(zeroTime)
active=YES(2019-08-15T15:21:59.756-07:00) health=(applnc:255
lm(t):5(2019-08-17T01:32:43.730-07:00) svc's)
  appliance id=6 address=10.0.0.6 lm(t):7(2019-08-08T03:50:55.470-07:00) tep
address=10.0.0.0/16 lm(t):6(2019-07-23T19:39:41.972-07:00) routable address=0.0.0.0
lm(t):0(zeroTime) oob address=172.31.170.230/21 lm(t):1(2019-08-18T23:14:28.727-07:00)
version=4.2(0.235j) lm(t):6(2019-08-15T15:22:00.562-07:00)
chassisId=066c943a-adbc-11e9-bbed-257398025731 lm(t):6(2019-08-15T15:22:00.562-07:00)
capabilities=0X3EFFFFFFFF--0X2020--0X20 lm(t):6(2019-08-14T07:55:20.053-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.820-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.821-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.821-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.821-07:00)
cntrlSbst=(APPROVED, WZP22350JFT) lm(t):6(2019-08-15T15:22:00.562-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=9
lm(t):6(2019-08-08T03:28:11.246-07:00) commissioned=YES lm(t):5(2019-08-08T03:03:57.387-07:00)
  registered=YES lm(t):7(2019-07-24T15:24:25.693-07:00) standby=NO
lm(t):6(2019-08-08T03:28:11.246-07:00) DRR=NO lm(t):5(2019-08-08T03:03:57.387-07:00) apicX=NO
  lm(t):6(2019-08-08T03:28:11.246-07:00) virtual=NO lm(t):0(zeroTime)
active=YES(2019-08-13T17:30:37.663-07:00) health=(applnc:255
lm(t):6(2019-08-15T15:57:05.128-07:00) svc's)
  appliance id=7 address=10.0.0.7 lm(t):7(2019-08-08T03:50:48.149-07:00) tep
address=10.0.0.0/16 lm(t):7(2019-07-24T15:24:19.988-07:00) routable address=0.0.0.0
lm(t):0(zeroTime) oob address=172.31.172.157/21 lm(t):1(2019-08-18T23:14:28.722-07:00)
version=4.2(0.235j) lm(t):7(2019-08-15T15:22:00.539-07:00)
chassisId=859be4ae-ae61-11e9-9840-7d9d67698989 lm(t):7(2019-08-15T15:22:00.539-07:00)
capabilities=0X3EFFFFFFFF--0X2020--0X40 lm(t):7(2019-08-14T07:55:23.872-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
cntrlSbst=(APPROVED, FCH2051V116) lm(t):7(2019-08-15T15:22:00.539-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=10
lm(t):7(2019-08-08T03:50:48.149-07:00) commissioned=YES lm(t):6(2019-08-08T03:28:16.727-07:00)
  registered=YES lm(t):6(2019-07-24T15:27:25.518-07:00) standby=NO
lm(t):7(2019-08-08T03:50:48.149-07:00) DRR=NO lm(t):6(2019-08-08T03:28:16.727-07:00) apicX=NO
  lm(t):7(2019-08-08T03:50:48.149-07:00) virtual=NO lm(t):0(zeroTime)
active=YES(2019-08-13T17:30:45.488-07:00) health=(applnc:255
lm(t):7(2019-08-17T01:39:26.549-07:00) svc's)
-----
clusterTime=<diff=2817 common=2019-08-19T15:33:55.929-07:00
local=2019-08-19T15:33:53.112-07:00 pF=<displForm=0 offsSt=0 offsVlu=-25200
lm(t):7(2019-08-08T03:50:55.925-07:00)>>
-----

```

```
apicl# acidiag rvread 6 3 1
```

```
(6,3,1) st:6 lm(t):3(2014-10-16T08:48:20.238+00:00) le: reSt:LEADER voGr:0 cuTerm:0x19
lCoTe:0x18
```

```

lCoIn:0x1800000000001b2a veFiSt:0x31 veFiEn:0x31 lm(t):3(2014-10-16T08:48:20.120+00:00)

lastUpdt 2014-10-16T09:07:00.214+00:00
-----
clusterTime=<diff=65247252 common=2014-10-16T09:07:01.837+00:00
local=2014-10-15T14:59:34.585+00:00
pF=<displForm=0 offsSt=0 offsVlu=0 lm(t):3(2014-10-16T04:50:08.714+00:00)>>

apic1# acidiag rvread 6 3
(6,3,1) st:6 lm(t):3(2014-10-16T08:48:20.238+00:00) le: reSt:LEADER voGr:0 cuTerm:0x19
lCoTe:0x18
lCoIn:0x1800000000001b2a veFiSt:0x31 veFiEn:0x31 lm(t):3(2014-10-16T08:48:20.120+00:00)

lastUpdt 2014-10-16T09:08:30.240+00:00
(6,3,2) st:6 lm(t):1(2014-10-16T08:47:25.323+00:00) le: reSt:FOLLOWER voGr:0 cuTerm:0x19
lCoTe:0x18
lCoIn:0x1800000000001b2a veFiSt:0x49 veFiEn:0x49 lm(t):1(2014-10-16T08:48:20.384+00:00)
lp: clSt:2
lm(t):1(2014-10-16T08:47:03.286+00:00) dbSt:2 lm(t):1(2014-10-16T08:47:02.143+00:00)
stMmt:1
lm(t):0(zeroTime) dbCrTs:2014-10-16T08:47:02.143+00:00 lastUpdt
2014-10-16T08:48:20.384+00:00
(6,3,3) st:6 lm(t):2(2014-10-16T08:47:13.576+00:00) le: reSt:FOLLOWER voGr:0 cuTerm:0x19
lCoTe:0x18
lCoIn:0x1800000000001b2a veFiSt:0x43 veFiEn:0x43 lm(t):2(2014-10-16T08:48:20.376+00:00)

lastUpdt 2014-10-16T09:08:30.240+00:00
-----
clusterTime=<diff=65247251 common=2014-10-16T09:08:30.445+00:00
local=2014-10-15T15:01:03.194+00:00
pF=<displForm=0 offsSt=0 offsVlu=0 lm(t):3(2014-10-16T04:50:08.714+00:00)>>

```

