Cisco APIC Troubleshooting Guide

Last Modified: 2016-09-28

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

CHAPTER 15

CHAPTER 16

APPENDIX A

APPENDIX B

APPENDIX C

Preface

This preface includes the following sections:

Audience

This guide is intended for system and network engineers with a background in troubleshooting data systems, networks, and storage systems.

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.

Convention	Description
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
`variable`	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

⚠ **Warning** IMPORTANT SAFETY INSTRUCTIONS

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Related Documentation

Cisco Application Centric Infrastructure (ACI) Documentation

The ACI documentation is available at the following URL: http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html.

Cisco Application Centric Infrastructure (ACI) Simulator Documentation

The Cisco ACI Simulator documentation is available at http://www.cisco.com/c/en/us/support/cloud-systems-management/application-centric-infrastructure-simulator/tsd-products-support-series-home.html.

Cisco Nexus 9000 Series Switches Documentation

The Cisco Nexus 9000 Series Switches documentation is available at http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/tsd-products-support-series-home.html.

Cisco Application Virtual Switch Documentation

The Cisco Application Virtual Switch (AVS) documentation is available at http://www.cisco.com/c/en/us/support/switches/application-virtual-switch/tsd-products-support-series-home.html.

Cisco Application Centric Infrastructure (ACI) Integration with OpenStack Documentation

Cisco ACI integration with OpenStack documentation is available at http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to apic-docfeedback@cisco.com. We appreciate your feedback.

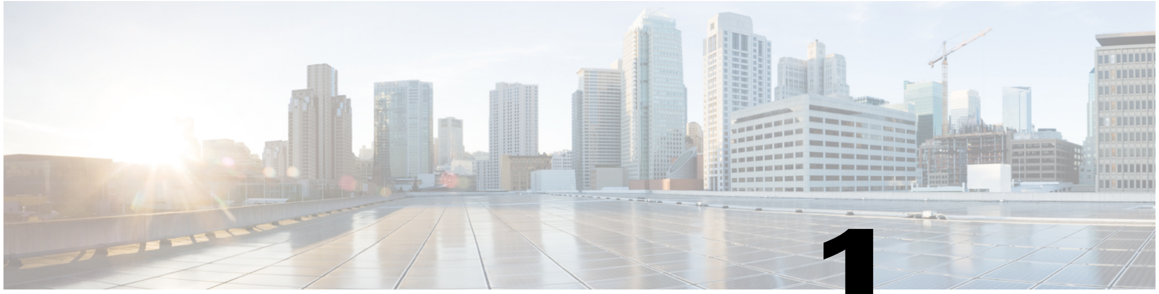

CHAPTER 1

New and Changed

New and Changed Information

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

Cisco APIC Release Version	Feature	Description
Release 2.1(1h)	Troubleshooting EVPN Type-2 Route Advertisement	Troubleshooting EVPN Type-2 Route Distribution to a DCIG
	Troubleshooting QoS policies	Troubleshooting QoS Policies, on page 133
	Troubleshooting Unwanted UI Objects	Removing Unwanted _ui_ Objects

Cisco APIC Release Version	Feature	Description
Release 2.0(1f)	Enabling and Viewing Digital Optical Monitoring Statistics	Enabling and Viewing Digital Optical Monitoring Statistics, on page 32
	Enabling and Viewing ACL Contract Permit and Deny Logs	Enabling and Viewing ACL Contract Permit and Deny Logs
	Syslog in NX-OS-Style CLI Format	Enabling Syslog to Display in NX-OS CLI Format, Using the REST API, on page 54
	Port Security	Confirming the Port Security Installation, on page 129
	Layer 3 multicast with multicast routing enabled using the Protocol Independent Multicast (PIM) protocol	Determining Why a PIM Interface Was Not Created, on page 127
Release 1.2(2)	Port Tracking Policy for Uplink Failure Detection	Enabling Port Tracking for Uplink Failure Detection, on page 37

CHAPTER 2

Troubleshooting Overview

The chapters in this guide describe common troubleshooting tips for specific Cisco APIC features and provide information about monitoring tools you can use for troubleshooting problems.

The features, issues, and tasks covered in this guide are listed below.

- **_ui_ Objects**—Explains how to remove unwanted _ui_ objects caused by making changes with the **Basic Mode** or the NX-OS CLI before using the **Advanced Mode**.
- **acidiag**—Explains how to use the acidiag command for troubleshooting operations on the Cisco APIC.
- **Cisco APIC Cluster**—Explains how to diagnose cluster faults and troubleshoot common cluster issues.
- **Cisco APIC Password Recovery and Emergency/Hidden Login Access**—Explains how to recover a password, how to access the rescue-user login to run troubleshooting commands, including erasing the configuration, and how to access a hidden login domain in case of a lockout.
- **Cisco APIC Troubleshooting Operations**—Explains how to gather information about your switches and how perform troubleshooting operations such as shutting down the system, shutting down the Cisco APIC controller, reloading the APIC controller, and turning on the LED locator.
- **Cisco APIC Troubleshooting Tools**—Explains how to use the Cisco APIC troubleshooting tools for monitoring traffic, debugging, and detecting issues such as traffic drops, misrouting, blocked paths, and uplink failures.
- **Endpoint Connectivity**—Explains how to troubleshoot endpoint connectivity using the Cisco APIC troubleshooting tools, such as traceroute, atomic counters, and SPAN, and how to connect an SFP module to a new card.

  Note Information about the Cisco APIC troubleshooting tools is located in the Using the Cisco APIC Troubleshooting Tools, on page 19 chapter.

- **EVPN Type-2 Host Routes**—Provides verification steps for this feature.
- **Export Policies**—Enables you to export statistics, tech support collections, faults and events, and to process core files and debug data from the fabric to any external hosts in a variety of formats.
- **Fabric Rebuild**—Explains how to rebuild your fabric.
- **IP-Based EPG**—Explains how to verify that you have correctly configured an IP-based EPG using the Cisco APIC GUI and using switch commands.

- **Leaf Connectivity—**Explains how to recover a disconnected leaf using the REST API.
- **PIM Interfaces—**Explains what to check when a PIM interface is not created for an L3Out, a multicast tunnel interface, or for a multicast-enabled bridge domain.
- **Port Security—**Explains how to confirm your port security hardware and software installations.
- **QoS—**Provides specific troubleshooting scenarios for this feature.
- **SSL Ciphers—**Explains how to determine if an SSL cipher is supported.
- **Switch Inventory—**Explains how to find the switch serial and model numbers. This helps TAC troubleshoot issues that you may experience.

Troubleshooting Basics

The following are basic steps for troubleshooting:

Before You Begin

- Familiarize yourself with the tools listed in Using the Cisco APIC Troubleshooting Tools, on page 19.
- Familiarize yourself with the Cisco APIC Troubleshooting Operations, on page 15.
- For issues with a specific feature, check the main contents of this guide for your feature. Troubleshooting tips are listed per-feature.

Step 1 Gather information that defines the specific symptoms.

Note In many cases, you can use the tools listed and described in the Using the Cisco APIC Troubleshooting Tools, on page 19 chapter to gather useful troubleshooting information.

Step 2 Identify all potential problems that could be causing the symptoms.

Step 3 Systematically eliminate each potential problem (from most likely to least likely) until the symptoms disappear.

Note This guide provides step-by-step instructions for confirming the installation and configuration of specific features such as port security, endpoint connectivity, PIM, and IP-based EPGs. Following the instructions can help you to narrow down and resolve problems you are experiencing.

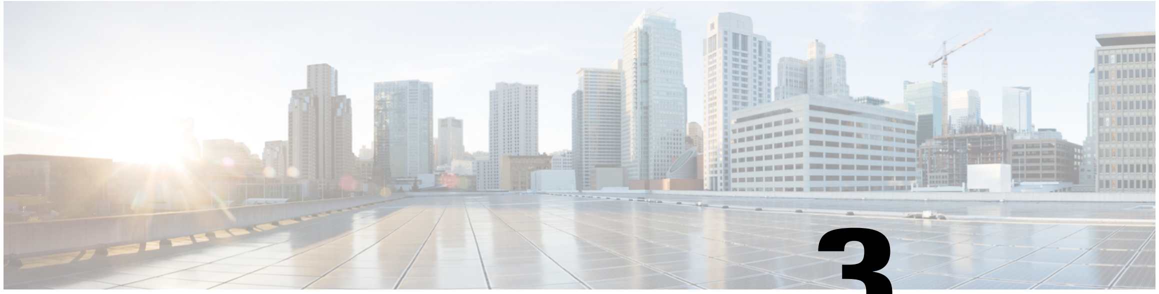

CHAPTER 3

Troubleshooting the Cisco APIC Cluster

This chapter contains information about cluster faults and possible solutions for common scenarios such as when a cluster does not expand and when a Cisco APIC has any of the following issues:

- Fails in the cluster
- Loses connection to the fabric and leaf switch
- Cannot communicate with another Cisco APIC after a reboot
- Joins a cluster when decommissioned
- Displays faults during changes to the cluster size
- Is unable to join a cluster
- Is unreachable in a cluster
- Is down
- Has a mismatched ChassisID after a reboot

This chapter contains the following sections:

Cluster Troubleshooting Scenarios

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

Problem	Solution
An APIC node fails within the cluster. For example, node 2 of a cluster of 5 APICs fails.	There are two available solutions: • Leave the target size and replace the APIC. • Reduce the cluster size to 4, decommission controller 5, and recommission it as APIC 2.The target size remains 4, and the operational size is 4 when the reconfigured APIC becomes active. **Note** You can add a replacement APIC to the cluster and expand the target and operational size. For instructions on how to add a new APIC, refer to the *Cisco APIC Management, Installation, Upgrade, and Downgrade Guide*..
A new APIC connects to the fabric and loses connection to a leaf switch.	Use the following commands to check for an infra (infrastructure) VLAN mismatch: • cat /mit/sys/lldp/inst/if-\[eth1--1\]/ctrlradj/summary**—**Displays the VLAN configured on the leaf switch. • cat /mit/sys/lldp/inst/if-\[eth1--1\]/ctrlradj/summary**—**Displays the infra (infrastructure) VLANs advertised by connected APICs. If the output of these commands shows different VLANs, the new APIC is not configured with the correct infra (infrastructure) VLAN. To correct this issue, follow these steps: • Log in to the APIC using rescue-user. **Note** Admin credentials do not work because the APIC is not part of the fabric. • Erase the configuration and reboot the APIC using the **acidiag touch setup** command. • Reconfigure the APIC. Verify that the fabric name, TEP addresses, and infra (infrastructure) VLAN match the APICs in the cluster. • Reload the leaf node.
Two APICs cannot communicate after a reboot.	The issue can occur after the following sequence of events: • APIC1 and APIC2 discover each other. • APIC1 reboots and becomes active with a new ChassisID (APIC1a) • The two APICs no longer communicate. In this scenario, APIC1a discovers APIC2, but APIC2 is unavailable because it is in a cluster with APIC1, which appears to be offline. As a result, APIC1a does not accept messages from APIC2. To resolve the issue, decommission APIC1 on APIC2, and commission APIC1 again.

Problem	Solution
A decommissioned APIC joins a cluster.	The issue can occur after the following sequence of events: • A member of the cluster becomes unavailable or the cluster splits. • An APIC is decommissioned. • After the cluster recovers, the decommissioned APIC is automatically commissioned. To resolve the issue, decommission the APIC after the cluster recovers.
Mismatched ChassisID following reboot.	The issue occurs when an APIC boots with a ChassisID different from the ChassisID registered in the cluster. As a result, messages from this APIC are discarded. To resolve the issue, ensure that you decommission the APIC before rebooting.
The APIC displays faults during changes to cluster size.	A variety of conditions can prevent a cluster from extending the OperationalClusterSize to meet the AdminstrativeClusterSize. For more information, inspect the fault and review Cluster Faults.
An APIC is unable to join a cluster.	The issue occurs when two APICs are configured with the same ClusterID when a cluster expands. As a result, one of the two APICs cannot join the cluster and displays an expansion-contender-chassis-id-mismatch fault. To resolve the issue, configure the APIC outside the cluster with a new cluster ID.
APIC unreachable in cluster.	Check the following settings to diagnose the issue: • Verify that fabric discovery is complete. • Identify the switch that is missing from the fabric. • Check whether the switch has requested and received an IP address from an APIC. • Verify that the switch has loaded a software image. • Verify how long the switch has been active. • Verify that all processes are running on the switch. For more information, see the acidiag Command. • Confirm that the missing switch has the correct date and time. • Confirm that the switch can communicate with other APICs.

Problem	Solution
Cluster does not expand.	The issue occurs under the following circumstances: • The OperationalClusterSize is smaller than the number of APICs. • No expansion contender (for example, the admin size is 5 and there is not an APIC with a clusterID of 4. • There is no connectivity between the cluster and a new APIC • Heartbeat messages are rejected by the new APIC • System is not healthy • An unavailable appliance is carrying a data subset that is related to relocation • Service is down on an appliance with a data subset that is related to relocation • Unhealthy data subset related to relocation
An APIC is down.	Check the following: • Connectivity issue—Verify connectivity using ping. • Interface type mismatch—Confirm that all APICs are set to in-band communication. • Fabric connectivity—Confirm that fabric connectivity is normal and that fabric discovery is complete. • Heartbeat rejected—Check the fltInfraIICIMsgSrcOutsider fault. Common errors include operational cluster size, mismatched ChassisID, source ID outside of the operational cluster size, source not commissioned, and fabric domain mismatch.

Cluster Faults

The APIC supports a variety of faults to help diagnose cluster problems. The following sections describe the two major cluster fault types.

Discard Faults

The APIC discards cluster messages that are not from a current cluster peer or cluster expansion candidate. If the APIC discards a message, it raises a fault that contains the originating APIC's serial number, cluster ID, and a timestamp. The following table summarizes the faults for discarded messages:

Fault	Meaning
expansion-contender-chassis-id-mismatch	The ChassisID of the transmitting APIC does not match the ChassisID learned by the cluster for expansion.

Fault	Meaning
expansion-contender-fabric-domain-mismatch	The FabricID of the transmitting APIC does not match the FabricID learned by the cluster for expansion.
expansion-contender-id-is-not-next-to-oper-cluster-size	The transmitting APIC has an inappropriate cluster ID for expansion. The value should be one greater than the current OperationalClusterSize.
expansion-contender-message-is-not-heartbeat	The transmitting APIC does not transmit continuous heartbeat messages.
fabric-domain-mismatch	The FabricID of the transmitting APIC does not match the FabricID of the cluster.
operational-cluster-size-distance-cannot-be-bridged	The transmitting APIC has an OperationalClusterSize that is different from that of the receiving APIC by more than 1. The receiving APIC rejects the request.
source-chassis-id-mismatch	The ChassisID of the transmitting APIC does not match the ChassisID registered with the cluster.
source-cluster-id-illegal	The transmitting APIC has a clusterID value that is not permitted.
source-has-mismatched-target-chassis-id	The target ChassisID of the transmitting APIC does not match the Chassis ID of the receiving APIC.
source-id-is-outside-operational-cluster-size	The transmitting APIC has a cluster ID that is outside of the OperationalClusterSize for the cluster.
source-is-not-commissioned	The transmitting APIC has a cluster ID that is currently decommissioned in the cluster.

Cluster Change Faults

The following faults apply when there is an error during a change to the APIC cluster size.

Fault	Meaning
cluster-is-stuck-at-size-2	This fault is issued if the OperationalClusterSize remains at 2 for an extended period. To resolve the issue, restore the cluster target size.
most-right-appliance-remains-commissioned	The last APIC within a cluster is still in service, which prevents the cluster from shrinking.
no-expansion-contender	The cluster cannot detect an APIC with a higher cluster ID, preventing the cluster from expanding.
service-down-on-appliance-carrying-replica-related-to-relocation	The data subset to be relocated has a copy on a service that is experiencing a failure. Indicates that there are multiple such failures on the APIC.
unavailable-appliance-carrying-replica-related-to-relocation	The data subset to be relocated has a copy on an unavailable APIC. To resolve the fault, restore the unavailable APIC.

Fault	Meaning
unhealthy-replica-related-to-relocation	The data subset to be relocated has a copy on an APIC that is not healthy. To resolve the fault, determine the root cause of the failure.

APIC Unavailable

The following cluster faults can apply when an APIC is unavailable:

Fault	Meaning
fltInfraReplicaReplicaState	The cluster is unable to bring up a data subset.
fltInfraReplicaDatabaseState	Indicates a corruption in the data store service.
fltInfraServiceHealth	Indicates that a data subset is not fully functional.
fltInfraWiNodeHealth	Indicates that an APIC is not fully functional.

CHAPTER 4

Recovering Cisco APIC Passwords and Accessing Special Logins

This chapter explains how to recover your Cisco APIC password, how to access the rescue-user login to run troubleshooting commands, including the command for erasing the configuration, and how to access a hidden login domain that allows you to log in using the local user database in case of a lockout.

This chapter contains the following sections:

Recovering the APIC Password

Follow these steps to recover the APIC password.

Step 1 Create and save an empty file named "aci-admin-passwd-reset.txt".

Step 2 Add the file to a USB drive.

Step 3 Connect the USB drive to one of the rear USB ports on the Cisco APIC.

Step 4 Reboot the APIC using Cisco Integrated Management Controller (CIMC) or by hard power cycling the device.

Step 5 When the APIC displays the "Press any key to enter the menu" prompt, press a key to interrupt the boot process.

Step 6 The APIC displays supported Linux versions. Highlight the version installed on your system and press **e** to edit the boot command.

Step 7 Highlight the kernel and press **e** to edit the command in boot sequence.

Step 8 Add the name of the empty file to the end of the command, shown as follows:

Example:

```
[ Minimal BASH-like line editing is supported.  For the first word, TAB
  lists possible command completions. Anywhere else TAB lists the possible
  completions of a device/filename. ESC at any time cancels. ENTER
  at any time accepts your changes.]
< rhgb quiet selinux=0 audit=1  aci-admin-passwd-reset
```

Step 9 Press **Enter** to save the file.

Step 10 Press **b** to boot the APIC.

Note To cancel the password reset operation and return to the default boot parameters, press **Esc** and **Enter**.

Step 11 The APIC boots and prompts for a new administrator password.

Using the Rescue-user Account to Erase the Cisco APIC Configuration Using the NX-OS Style CLI

The rescue-user is an emergency login that provides access to the APIC even when it is not in a cluster. You can use this login to run troubleshooting commands including erasing the configuration.

Note If the APIC is part of a healthy cluster, the rescue-user account is protected with the admin password.

Step 1 Access the APIC using the Cisco Integrated Management Controller (CIMC) console.

Step 2 Login as rescue-user.

Note If an admin password is in place and the APIC is logged onto the fabric, the rescue-user password is the same as the admin password. Otherwise there is no rescue-user password.

Step 3 Use the **acidiag touch** command to clear the configuration.

Example:

```
apic1# acidiag touch setup
```

Using the Fallback Login Domain to Log in to the Local Database

There is a hidden login domain named "fallback" that allows you to log in using the local user database in case of lockout. The format of the username used for the authentication method is `apic#fallback\\<username>`.

Use the fallback login domain to log in to the local database in the GUI:

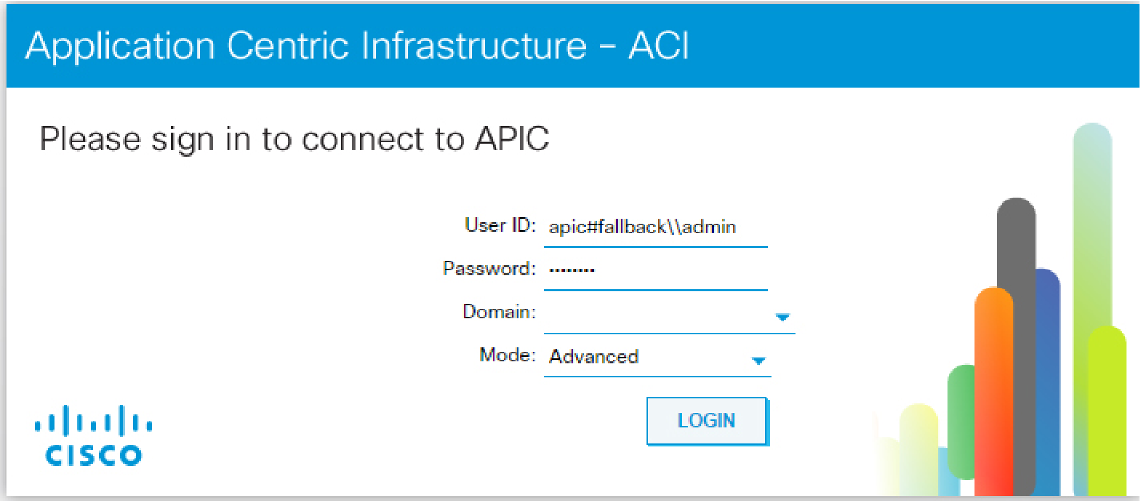
Application Centric Infrastructure – ACI
Please sign in to connect to APIC
LOGIN
cisco

User ID: apic#fallback\\admin

Password: ········

Domain:

Mode: Advanced

Or, log in to the fallback login domain using the NX-OS CLI, shown as follows:

```
apic1(config)# aaa authentication login domain fallback
apic1(config-domain)# ?
group Set provider group for login domain
realm Specify server realm
```

Or, you can use the REST API to log in to the fallback login domain, shown as follows:

- URL: https://ifav41-ifc1/api/aaaLogin.xml
- DATA:

```
<aaaUser name="apic#fallback\\admin"
pwd="passwordhere"/>
```

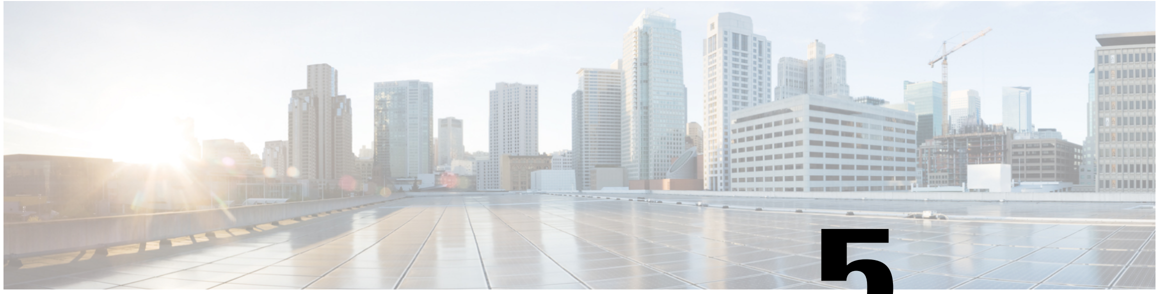

CHAPTER 5

Cisco APIC Troubleshooting Operations

This chapter explains how to perform the basic troubleshooting operations and contains the following sections:

Shutting Down the APIC System

This procedure explains how to shut down the APIC system.

Note After you shut down the system, you will move it (re-locate the entire fabric) then power it up, then update the time zone and/or NTP servers accordingly.

1 Shut down one APIC at a time by right-clicking on it then selecting **Shutdown** from the pull-down menu:

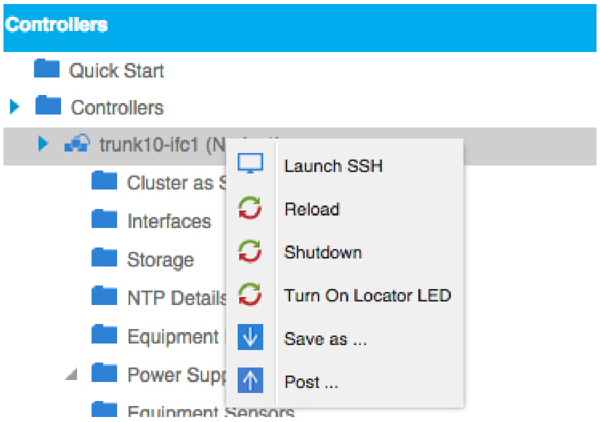

2 Start up the APIC at the new location.

3 Check that the cluster has fully converged as follows:

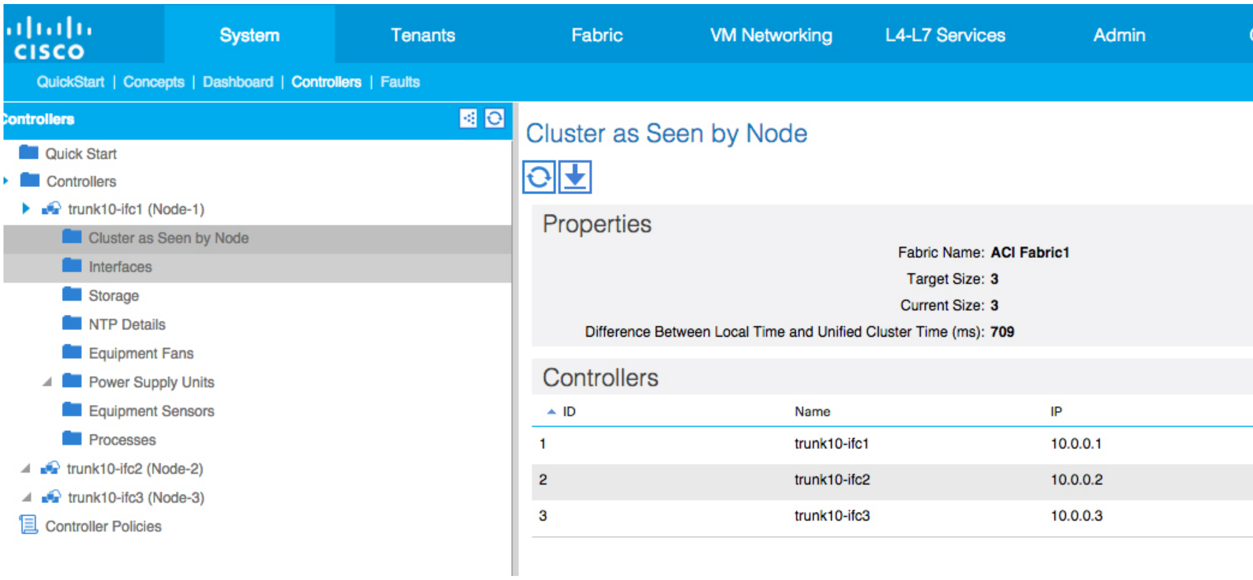

Cluster as Seen by Node

Properties

Fabric Name: **ACI Fabric1**
Target Size: **3**
Current Size: **3**
Difference Between Local Time and Unified Cluster Time (ms): **709**

Controllers

ID	Name	IP
1	trunk10-ifc1	10.0.0.1
2	trunk10-ifc2	10.0.0.2
3	trunk10-ifc3	10.0.0.3

4 Proceed with the next APIC.

Before You Begin

Ensure cluster health is fully fit.

Shutting Down the APIC Controller Using the GUI

This document describes how to shut down the APIC controller.

Note This procedure instructs on how to shut down the APIC controller only (not the entire APIC system itself). Following this procedure causes the controller to shut down immediately. Use caution in performing a shutdown because the only way to bring the controller back up is to do so from the actual machine. If you need to access the machine, refer to the "Turning on the Locator LED Using the GUI" section in this chapter.

Shut down a single APIC controller as follows:

Note If possible, move APICs one at a time. As long as there are at least two APICs in the cluster online, there is read/write access. If you need to relocate more than one APIC at a time, this results in one or no remaining controllers online, and the fabric will go into a read-only mode when they are shutdown. During this time there can be no policy changes including Endpoint moves (including Virtual Machine movement). Once the APICs are shut down using the following procedure, relocate the controller, and power it back up under the new rack. Then, confirm that the cluster health returns to fully fit status.

1 In the menu bar, click **System**.

2 In the submenu bar, click **Controllers**.

3 Under **Controllers,** click the APIC node that you would like to reload, for example, **apic1 (Node-1).**

4 In the right window pane, at the top of the screen, click the **General** tab.

5 In the right window pane, at the top of the screen and under the tabs, click the **ACTIONS** pull-down menu.

6 Select **Shutdown** from the pull-down menu to immediately reload the APIC controller.

Note Another way to use this Shutdown option is to right-click on the APIC node (such as **apic1 (Node-1)**, and select **Shutdown** from the pull-down list.

7 Relocate the controller, then power it up.

8 Confirm cluster health returns to fully fit status.

Using the APIC Reload Option Using the GUI

This document describes how to reload the APIC controller (not the entire APIC system) using the GUI.

Reload the APIC controller as follows:

1 In the menu bar, click **System**.

2 In the submenu bar, click **Controllers**.

3 Under Controllers, click the APIC node that you would like to reload, for example, **apic1 (Node-1).**

4 In the right window pane, at the top of the screen, click the **General** tab.

5 In the right window pane, at the top of the screen and under the tabs, click the **ACTIONS** pull-down menu.

6 Select **Reload** from the pull-down menu to immediately reload the APIC controller.

Note Another way to use this Reload option is to right-click on the APIC node (such as **apic1 (Node-1)**, and select **Reload** from the pull-down list.

Controlling the LED Locator Using the GUI

This document describes how to turn on the LED locator for the APIC controller using the GUI.

Turn on (or turn off) the LED locator of the APIC controller using the GUI as follows:

1 In the menu bar, click **System**.

2 In the submenu bar, click **Controllers**.

3 Under Controllers, click the APIC node that you would like to reload, for example, **apic1 (Node-1).**

4 In the right window pane, at the top of the screen, click the **General** tab.

5 In the right window pane, at the top of the screen and under the tabs, click the **ACTIONS** pull-down menu.

6 Select **Turn On LED Locator** (or **Turn Off LED Locator**) from the pull-down menu.

Note Another way to use this option is to right-click on the APIC node (such as **apic1 (Node-1)**, and select **Turn On LED Locator** (or **Turn Off LED Locator**) from the pull-down list.

CHAPTER 6

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 Analizer (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.

This chapter contains the following sections:

Enabling and Viewing ACL Contract Permit and Deny Logs

To log and/or monitor the traffic flow for a contract rule, you can enable and view the logging of packets or flows that were allowed to be sent because of contract permit rules and the logging of packets or flows that were dropped because of taboo contract deny rules.

For information on contracts, see the *Cisco Application Centric Infrastructure Fundamentals* guide.

Enabling ACL Contract Permit Logging Using the GUI

The following steps show how to enable Contract permit logging using the GUI:

Step 1 On the menu bar, choose **Tenants** > **<tenant name>**.

Step 2 In the **Navigation** pane, expand **Security Policies**.

Step 3 Right-click **Contracts** and choose **Create Contract**.

Step 4 In the Create Contract dialog box, perform the following actions:

a) In the **Name** field, type the name for the contract.
b) In the **Scope** field, choose the scope for it (VRF, Tenant, or Global).
c) Optional. Set the target DSCP or QoS class to be applied to the contract.

d) Expand **Subjects**.

Step 5 In the Create Contract Subject dialog box, perform the following actions:

Step 6 Type the name of the subject and an optional description.

Step 7 Optional. From the drop-down list for the target DSCP, select the DSCP to be applied to the subject.

Step 8 Leave **Apply Both Directions** checked, unless you want the contract to only be applied from the consumer to the provider, instead of in both directions.

Step 9 Leave **Reverse Filter Ports** checked if you unchecked **Apply Both Directions**, to swap the Layer 4 source and destination ports, so that the rule will be applied from the provider to the consumer.

Step 10 Expand **Filter Chain**.

Step 11 In the **Name** drop-down list, choose an option; for example, click **arp**, **default**, **est**, or **icmp**.

Step 12 In the **Directives** drop-down list, click **log**.

Step 13 Click **Update**.

Step 14 Click **OK**.

Step 15 Click **SUBMIT**.
Logging is enabled for this contract.

Enabling ACL Contract Permit Logging Using the NX-OS CLI

The following example shows how to enable Contract permit logging using the NX-OS CLI.

Step 1 To enable logging of packets or flows that were allowed to be sent because of Contract permit rules, use the following commands:

```
configure
tenant <tenantName>
contract <contractName> type <permit>
subject <subject Name>
access-group <access-list> <in/out/both> log
```

Example:

For example:

```
apic1# configure
apic1(config)# tenant BDMode1
apic1(config-tenant)# contract Logicmp type permit
apic1(config-tenant-contract)# subject icmp
apic1(config-tenant-contract-subj)# access-group arp both log
```

Step 2 To disable the permit logging use the **no** form of the access-group command; for example, use the `no access-group arp both log` command.

Enabling ACL Contract Permit Logging Using the REST API

The following example shows you how to enable Contract permit logging using the REST API.

To enable Contract permit logging, post data such as the following example:

```
POST https://192.0.20.123/api/node/mo/uni/tn-sgladwin_t1/brc-ICMP_Contract.json
{
"vzBrCP":{
   "attributes":{
      "dn" : "uni/tn-sgladwin_t1/brc-ICMP_Contract",
      "name" : "ICMP_Contract",
      "rn" : "brc-ICMP_Contract","
      "status" : "created"},
   "children":[{
      "vzSubj":{
         "attributes":{
            "dn" : "uni/tn-sgladwin_t1/brc-ICMP_Contract/subj-Permit_Contract_ICMP",
            "name" : "Permit_Contract_ICMP","rn":"subj-Permit_Contract_ICMP",
            "status":"created"},
   "children":[{
      "vzRsSubjFiltAtt":{
            "attributes":{
            "status":"created,modified",
            "tnVzFilterName":"icmp",
            "directives":"log"},
   "children":[]}}]}}]}}
response: {"totalCount":"0","imdata":[]}
```

Enabling Taboo Contract Deny Logging Using the GUI

The following steps show how to enable Taboo Contract deny logging using the GUI:

Step 1 On the menu bar, choose **Tenants** > **<tenant name>**.

Step 2 In the **Navigation** pane, expand **Security Policies**.

Step 3 Right-click **Taboo Contracts** and choose **Create Taboo Contract**.

Step 4 In the Create Taboo Contract dialog box, perform the following actions to specify the Taboo contract:

a) In the **Name** field, type the name for the contract.
b) Optional. In the **Description** field, type a description of the Taboo contract.
c) Expand **Subjects**.

Step 5 In the Create Taboo Contract Subject dialog box, perform the following actions:

a) In the Specify Identity of Subject area, type a name and optional description.
b) Expand **Filters**.

c) From the **Name** drop-down list, choose one of the following values:**<tenant_name>/arp**, **<tenant_name>/default**, **<tenant_name>/est**, **<tenant_name>/icmp**, or **Create Filter.**

Note If you chose **Create Filter**, in the Specify Filter Identity Area, perform the following actions to specify criteria for the ACL Deny rule:

1. Type a name and optional description.
2. Expand **Entries**, type a name for the rule, and choose the ACL Deny criteria for the rule.
3. Click **Update**.
4. Click **Submit**.

Step 6 In the Directives drop-down list, click **log**.

Step 7 Click **Update**.

Step 8 Click **OK**.

Step 9 Click **SUBMIT**.
Logging is enabled for this Taboo contract.

Enabling Taboo Contract Deny Logging Using the NX-OS CLI

The following example shows how to enable Taboo Contract deny logging using the NX-OS CLI.

Step 1 To enable logging of packets or flows dropped because of Taboo Contract deny rules, use the following commands:

```
configure
tenant <tenantName>
contract <contractName> type <deny>
subject <subject Name>
access-group <access-list> <both> log
```

Example:
For example:

```
apic1# configure
apic1(config)# tenant BDMode1
apic1(config-tenant)# contract dropFTP type deny
apic1(config-tenant-contract)# subject dropftp
apic1(config-tenant-contract-subj)# access-group https both log
```

Step 2 To disable the deny logging use the **no** form of the access-group command; for example, use the `no access-group https both log` command.

Enabling Taboo Contract Deny Logging Using the REST API

The following example shows you how to enable Taboo Contract deny logging using the REST API.

To enable ACL deny logging, post data such as the follow example:

```
POST https://192.0.20.123/api/node/mo/uni/tn-sgladwin_t1/taboo-TCP_Taboo_Contract.json
{
"vzTaboo":{
   "attributes":{
      "dn":"uni/tn-sgladwin_t1/taboo-TCP_Taboo_Contract",
      "name":"TCP_Taboo_Contract",
      "rn":"taboo-TCP_Taboo_Contract",
      "status":"created"},
   "children":[{
      "vzTSubj":{
         "attributes":{
            "dn":"uni/tn-sgladwin_t1/taboo-TCP_Taboo_Contract/tsubj-TCP_Filter_Subject",
            "name":"TCP_Filter_Subject",
            "rn":"tsubj-TCP_Filter_Subject",
            "status":"created"},
         "children":[{
            "vzRsDenyRule":{
               "attributes":{
                  "tnVzFilterName":"TCP_Filter",
                  "directives":"log",
                  "status":"created"},
         "children":[]}}]}}]}}
response: {"totalCount":"0","imdata":[]}
```

Viewing ACL Permit and Deny Logs Using the GUI

The following steps show how to view ACL permit and deny logs (if they are enabled) for traffic flows, using the GUI:

Step 1 On the menu bar, choose **Tenants** > **<tenant name>**.

Step 2 In the **Navigation** pane, click on **Tenant <tenant name>**.

Step 3 In the **Tenants <tenant name> Work** pane, click the **Operational** tab.

Step 4 Under the **Operational** tab, click the **Flows** tab.
Under the **Flows** tab, click one of the tabs to view log data for Layer 2 permit logs (**L2 Permit**) Layer 3 permit logs (**L3 Permit**, Layer 2 deny logs (**L2 Drop**), or Layer 3 deny logs (**L3 Drop**). On each tab, you can view ACL logging data, if traffic is flowing. The data points differ according to the log type and ACL rule; for example, the following data points are included for **L3 Permit** logs:

- Timestamp

- VRF
- Source IP address
- Destination IP address
- Protocol
- Source Port
- Destination Port
- Source MAC address
- Destination MAC address
- Node (switch where data came from)
- Source interface
- VLAN
- VRF encapsulation

Note You can also use the **Packets** tab (next to the **Flows** tab) to access ACL logs for groups of packets (up to 10) with the same signature, source and destination. You can see what type of packets are being sent and which are being dropped.

Viewing ACL Permit and Deny Logs Using the REST API

The following example shows how to view permit and deny log data for traffic flows, using the REST API:

Before You Begin

You must enable permit or deny logging, before you can view ACL contract permit and deny log data.

Send the following query using the REST API:

```
GET
https://apic-ip-address/api/node/mo/uni/tn-sgladwin_t1.json?rsp-subtree-include=stats&rsp-subtree-class=fvOverallHealthHist15min
{
"totalCount":"1",
"imdata":[{
"fvTenant":{
   "attributes":{
       "childAction":"",
       "descr":"",
       "dn":"uni/tn-sgladwin_t1",
       "lcOwn":"local",
       "modTs":"2016-06-22T15:46:30.745+00:00",
       "monPolDn":"uni/tn-common/monepg-default",
       "name":"sgladwin_t1",
       "ownerKey":"",
        "ownerTag":"",
        "status":"",
        "uid":"15374"
}}}]}
```

Viewing ACL Permit and Deny Logs Using the NX-OS CLI

The following steps show how to view ACL log details using the NX-OS CLI **show acllog** command.

The full syntax for the command is **show acllog {permit | drop} l3 {pkt | flow} tenant <tenant name> vrf <vrf name> srcip <source ip> dstip <destination ip> srcport <source port> dstport <destination port> protocol <protocol> srcintf <source interface> start-time <startTime> end-time <endTime>**

Step 1 The following example shows how to use the **show acllog permit l3 pkt tenant <tenant name> vrf <vrf name> [detail]** command to display detailed information about the common VRF ACL Layer 3 permit packets that were sent:

```
apic1# show acllog permit l3 pkt tenant common vrf default detail
acllog permit l3 packets detail:
srcIp      : 10.2.0.19
dstIp      : 10.2.0.16
protocol   : udp
srcPort    : 13124
dstPort    : 4386
srcIntf    : port-channel5
vrfEncap   : VXLAN: 2097153
pktLen     : 112
srcMacAddr : 00:00:15:00:00:28
dstMacAddr : 00:00:12:00:00:25
timeStamp  : 2015-03-17T21:31:14.383+00:00
```

Step 2 The following example shows how to use the **show acllog** command to display information about common VRF Layer 3 UDP packets with source IP address 10.2.0.19, destination IP address 10.2.0.16, source port 13124, destination port 4386, source interface port-channel15, start -time 2015-03-17T21:00:00, and end-time 2015-03-18T00:00:00:

```
apic1# show acllog permit l3 pkt tenant common vrf copy srcip 10.2.0.19 dstip 10.2.0.16 srcport 13124
 dstport 4386
protocol 17 srcintf port-channel5 start-time 2015-03-17T21:00:00 end-time  2015-03-18T00:00:00
acllog Permit L3 Packets
     srcIp       dstIp         protocol   srcport   dstport   Node      srcIntf         vrfEncap
    pktLen   timeStamp
 ------------    -----------   -----     --------  -------   ------   ---------        -------------
   --------  --------------
   10.2.0.19     10.2.0.16      udp      13124      4386      101      port-channel5   VXLAN: 2097153
   112      2015-03-17T21:


                    31:14.383+00:00
```

Step 3 The following example shows how to use the **show acllog permit l2 pkt tenant <tenant name> vrf <vrf name> [detail]** command to view detailed information about default VRF Layer 2 packets that were sent:

```
apic1# show acllog permit l2 pkt tenant common vrf default detail

acllog permit l2 packets detail:
srcIntf    : port-channel5
pktLen     : 1
srcMacAddr : 00:00:66:00:00:66
dstMacAddr : 00:00:89:00:00:00
timeStamp  : 2015-03-17T21:31:14.383+00:00
```

Step 4 The following example shows how to use the **show acllog permit l2 pkt tenant <tenant name> vrf <vrf name> srcintf <s interface>** command to view information about default VRF Layer 2 packets sent from interface port-channel15:

```
apic1# show acllog permit l2 pkt tenant common vrf default srcintf port-channel5
acllog permit L2 Packets
      Node          srcIntf       pktLen     timeStamp
 --------------  --------------  --------  --------------
                 port-channel5      1      2015-03-17T21:
                                           31:14.383+00:00
```

Step 5 The following example shows how to use the **show acllog drop l3 pkt tenant <tenant name> vrf <vrf name> [detail]** command to show detailed information about packets that were dropped due to an ACL deny rule:

```
apic1# show acllog drop l3 pkt tenant common vrf copy detail
acllog drop l3 packets detail:
srcIp      : 10.2.0.19
dstIp      : 10.2.0.16
protocol   : udp
srcPort    : 13124
dstPort    : 4386
srcIntf    : port-channel5
vrfEncap   : VXLAN: 2097153
pktLen     : 112
srcMacAddr : 00:00:15:00:00:28
dstMacAddr : 00:00:12:00:00:25
timeStamp  : 2015-03-17T21:31:14.383+00:00
```

Step 6 The following example shows how to use the **show acllog drop l3 flow tenant <tenant name> vrf <vrf name> srcip <source ip> dstip <dst ip> srcport <src port> dstport <dst port> protocol <protocol> srcintf <src intf>** command to show information about UDP packets that were dropped, that had source IP address 10.2.0.19, destination IP address 10.2.0.16, source port 13124, destination port 4386, and souce interface port-channel15:

```
apic1# show acllog drop l3 pkt tenant common vrf copy srcip 10.2.0.19 dstip 10.2.0.16 srcPort 13124
 dstPort 4386 protocol 17 srcintf port-channel5
acllog drop L3 Packets
     srcIp           dstIp       protocol  srcPort   dstPort        Node           srcIntf
vrfEncap      pktLen     timeStamp
 --------------  --------------  --------  --------  --------  --------------  --------------
--------------  --------  --------------
   10.2.0.19       10.2.0.16       udp      13124      4386                   port-channel5   VXLAN:
 2097153    112     2015-03-17T21: 31:14.383+00:00
```

Step 7 The following example shows how to use the **show acllog drop l2 pkt tenant <tenant name> vrf <vrf name> [detail]** command to view detailed information about packets that were dropped because of an ACL deny rule:

```
apic1# show acllog drop l2 pkt tenant common vrf copy detail

acllog drop l2 packets detail:
srcIntf    : port-channel87
pktLen     : 1122
srcMacAddr : 00:00:11:00:00:11
dstMacAddr : 11:00:32:00:00:33
timeStamp  : 2015-03-17T21:31:14.383+00:00
```

Step 8 The following example shows how to use the **show acllog drop l2 pkt tenant <tenant name> vrf <vrf name> srcintf <srcintf name>** command to view information about packets dropped that originated from a specific interface:

```
apic1# show acllog drop l2 pkt tenant common vrf copy srcintf port-channel87
acllog drop L2 Packets
      Node          srcIntf       pktLen     timeStamp
 --------------  --------------  --------  --------------
                 port-channel87    1122    2015-03-17T21:
                                           31:14.383+00:00
```

Using Atomic Counter Policies for Gathering Statistics

Atomic counter policies enable you to gather statistics about your traffic between a combination of endpoints, endpoint groups, external interfaces, and IP addresses. The information gathered enables you to detect drops and misrouting in the fabric, which enables you to perform quick debugging and to isolate application connectivity issues.

Atomic Counters

Atomic Counters are useful for troubleshooting connectivity between endpoints, EPGs, or an application within the fabric. A user reporting application may be experiencing slowness, or atomic counters may be needed for monitoring any traffic loss between two endpoints. One capability provided by atomic counters is the ability to place a trouble ticket into a proactive monitoring mode, for example when the problem is intermittent, and not necessarily happening at the time the operator is actively working the ticket.

Atomic counters can help detect packet loss in the fabric and allow the quick isolation of the source of connectivity issues. Atomic counters require NTP to be enabled on the fabric.

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

- Counts of drops, admits, and excess packets
- 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

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.

Tenant atomic counters can provide the following:

- Application-specific counters for traffic across the fabric, including drops, admits, and excess packets
- Modes include the following:
- Endpoint to endpoint MAC address, or endpoint to endpoint IP address. Note that a single target endpoint could have multiple IP addresses associated with it.
- EPG to EPG with optional drill down
- EPG to endpoint
- EPG to * (any)
- Endpoint to external IP address

Note Atomic counters track the amount packets of between the two endpoints and use this as a measurement. They do not take into account drops or error counters in a hardware level.

Dropped packets are calculated when there are less packets received by the destination than transmitted by the source.

Excess packets are calculated when there are more packets received by the destination than transmitted by the source.

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 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 and/or destination counts only the traffic that is from/to the MAC and IP addresses that are present in the fvCEp managed objects (MOs). If the fvCEp MO has an empty IP address field, then all traffic to/from that MAC address would be counted regardless of the IP address. If the APIC has learned multiple IP addresses for an fvCEp, then traffic from only the one IP address in the fvCEp MO itself is counted as previously stated. In order to configure an atomic counter policy to/from a specific IP address, use the fvIp MO as the source and/or destination.
- If there is an fvIp behind an fvCEp, you must add fvIP-based policies and not fvCEp-based policies.

Configuring Atomic Counters

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 **Troubleshoot Policies**.

Step 4 Under **Troubleshoot Policies**, 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**)

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

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:

Step 1 On the menu bar, choose **Fabric** > **Fabric Policies**.

Step 2 In the **Navigation** pane, expand **Switch Policies** > **Policies** > **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:

a) In the **Name** field, enter a name for the policy.
b) Optional. In the **Description** field, enter a description of the policy.
c) Check 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:

a) In the **Name** field, enter a name for the policy group.
b) 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**.
c) Click **SUBMIT**.

Step 10 Attach the policy group you created to a switch as follows:

a) In the **Navigation** pane, expand **Switch Policies** > **Profiles**.
b) In the **Work** pane, click the **ACTIONS** drop-down menu and select **Create Leaf Switch Profile** or **Create Spine Switch Profile**, as appropriate.
c) In the dialog box, enter a name for the profile in the **Name** field.
d) Add the name of the switch you want associated with the profile under **Switch Associations**.
e) From the **Blocks** pull-down menu, check the boxes next to the applicable switches.

f) From the **Policy Group** pull-down menu, select the policy group you created earlier.
g) 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:

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.

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 statitistics 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.

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

You can view a summary of faults as follows:

- System Faults—Choose **System** > **Faults**.
- Tenant Faults-—Click **Tenants** > *tenant-name*, click the tenant name in the left pane, and choose the **Faults** tab in the right pane.
- Fabric Faults—Click **Fabric**, click a Pod in the left pane, and click the **Faults** tab in the right pane.

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 Uplink Failure Detection

Uplink failure detection can be enabled in the fabric access global port tracking policy. The port tracking policy monitors the status of links between leaf switches and spine 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.

Note In the advanced GUI, port tracking is located under **Fabric** > **Access Policies** > **Global Policies** > **Port Tracking**.

In the basic GUI, port tracking is located under **System** > **System Settings** > **Port Tracking**.

Depending on the model of leaf switch, each leaf switch can have 6, 8, or 12 uplink connections to each spine switch. The port tracking policy specifies the number of uplink connections that trigger the policy, and a delay timer for bringing the leaf switch access ports back up after the number of specified uplinks is exceeded.

The following example illustrates how a port tracking policy behaves:

- The leaf switches each have 6 active uplink connections to the spine switches.
- The port tracking policy specifies that the threshold of active uplink connections each leaf switch that triggers the policy is 2.
- The port tracking policy triggers when the number of active uplink connections from the leaf switch to the spine switches drops to 2.
- Each leaf switch monitors its uplink connections and triggers the port tracking policy according to the threshold specified in the policy.
- When the uplink 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 links that triggers port tracking is too high, all leaf switch access ports will be brought down.

Port Tracking Using the GUI

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

Step 1 From the **Fabric** menu, select **Access Policies**.

Step 2 In the left navigation pane of the **Access Policies** screen, select **Global Polices**.

Step 3 Under **Global Policies** tab, select the **Port Tracking** tab.

Step 4 Turn on the Port Tracking feature by selecting **on** next to **Port tracking state** under Properties.

Step 5 Turn off the Port Tracking feature by selecting **off** next to **Port tracking state** under Properties.

Step 6 Set the **Delay restore timer**, which is the configuration parameter used to specify the number of seconds before restoring and bringing up the downlinks after the fabric port tracking starts.

Step 7 Enter the maximum number of links (any configuration value from 0 - 12) that are left up before you trigger this configuration. (The default is zero.)

Step 8 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.

Step 1 Turn on the Port Tracking feature as follows:

Example:

```
apic1# 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:

```
apic1# 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.

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>
```

Configuring SNMP for Monitoring and Managing Devices

This section explains how to configure SNMP using the GUI.

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 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.

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

SNMP Access Support in ACI

SNMP support in ACI is as follows:

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

Note ACI supports a maximum of 10 trap receivers.

- SNMPv3 is supported by leaf and spine switches and by APIC.

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 APIC controller is added. Previous releases support SNMP only for leaf and spine switches.

For the complete list of MIBs supported in ACI, see http://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/mib/list/mib-support.html.

Configuring the SNMP Policy Using the GUI

This procedure configures and enables the SNMP policy on ACI 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 APIC out-of-band IP addresses in the 'mgmt' tenant. Although the out-of-band addresses are configured during APIC setup, the addresses must be explicitly configured in the 'mgmt' tenant before the out-of-band contract will take effect.

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

Step 7 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**.

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:

a) In the **Name** field, enter a pod policy group name.

b) 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.

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:

a) In the **Name** field, enter an SNMP destination name and click **Next**.

b) In the **Create Destinations** table, click the + icon to open the **Create SNMP Trap Destination** dialog box.

c) In the **Host Name/IP** field, enter an IP address or a fully qualified domain name for the destination host.

Note Cisco APIC Release 1.2(2) and later releases support IPv6 SNMP trap destinations.

d) Choose the **Port** number and SNMP **Version** for the destination.

e) For SNMP v1 or v2c destinations, enter one of the configured community names as **Security Name** and choose **noauth** as **v3 Security Level**.

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

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.

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:

a) In the **Name** field, enter an SNMP policy name.
b) 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.
c) 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.

Configuring SPAN for Traffic Monitoring

This section lists the SPAN guidelines and restrictions and explains how to configure SPAN sessions.

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.

Multinode SPAN

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

SPAN Guidelines and Restrictions

- Use SPAN for troubleshooting. SPAN traffic competes with user traffic for switch resources. To minimize the load, configure SPAN to copy only the specific traffic that you want to analyze.
- You cannot specify an l3extLIfP layer 3 subinterface as a SPAN source. You must use the entire port for monitoring traffic from external sources.
- Tenant and access SPAN use the encapsulated remote extension of SPAN (ERSPAN) type I, while fabric SPAN uses ERSPAN type II. 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 IPs must be learned in the fabric as an endpoint.
- SPAN supports IPv6 traffic but the destination IP for the ERSPAN cannot be an IPv6 address.
- See the *Verified Scalability Guide for Cisco ACI* document for SPAN-related limits, such as the maximum number of active SPAN sessions.

Configuring a SPAN Session

This procedure shows how to configure a SPAN policy to forward replicated source packets to a remote traffic analyzer.

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 **Troubleshooting Policies**, and expand **SPAN**.

Step 4 Under **SPAN**, right-click **SPAN Destination Groups** and choose **Create SPAN Destination Group**.

Step 5 In the **Create SPAN Destination Group** dialog box, perform the following actions:

a) In the **Name** field, enter a name for the SPAN destination group.
b) In the **Create Destinations** table, click the **+** icon to open the **Create SPAN Destination** dialog box.
c) In the **Name** field, enter a name for the SPAN destination.
d) From the **Destination EPG** drop-down lists, choose or enter the destination tenant, application profile, or EPG to which replicated packets will be forwarded.
e) In the **Destination IP** field, enter the IP address of the remote server that will receive the replicated packets.
f) In the **Source IP Prefix** field, enter the base IP address of the IP subnet of the source packets.
g) (Optional) In the **Flow ID** field, increment or decrement the flow ID value of the SPAN packet.
h) (Optional) In the **TTL** field, increment or decrement the IP time-to-live (TTL) value of the packets in the SPAN traffic.
i) (Optional) In the **MTU** field, increment or decrement the MTU truncation size for the packets.
j) (Optional) In the **DSCP** field, increment or decrement the IP DSCP value of the packets in the SPAN traffic.
k) Click **OK** to save the SPAN destination.
l) Click **Submit** to save the SPAN destination group.

Step 6 Under **SPAN**, right-click **SPAN Source Groups** and choose **Create SPAN Source Group**.

Step 7 In the **Create SPAN Source Group** dialog box, perform the following actions:

a) In the **Name** field, enter a name for the SPAN source group.
b) From the **Destination Group** drop-down list, choose the SPAN destination group that you configured previously.
c) In the **Create Sources** table, click the **+** icon to open the **Create ERSPAN Source** dialog box.
d) In the **Name** field, enter a name for the source.
e) In the **Direction** field, choose the radio button based on whether you want to replicate and forward packets that are incoming to the source, outgoing from the source, or both incoming and outgoing.
f) From the **Source EPG** drop-down list, choose the EPG (identified by Tenant/ApplicationProfile/EPG) whose packets will be replicated and forwarded to the SPAN destination.
g) Click **OK** to save the SPAN source.
h) Click **Submit** to save the SPAN source group.

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.

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:

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.

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 *Cisco Nexus 9000 Series NX-OS Configuration Guides*.
show platform internal counters port [*port_num* | **detail** | **nz** | {**internal** [**nz** | *int_port_num*]}]	Displays spine port statistics • *port_num*—Front port number without the slot. • **detail**—Returns SNMP, class and forwarding statistics. • **nz**—Displays only non-zero values. • **internal**—Displays internal port statistics. • *int_port_num*—Internal logical port number. For example, for BCM-0/97, enter 97. **Note** If there is a link reset, the counters will be zeroed out on the switch. The conditions of counter reset include the following: • accidental link reset • manually enabled port (after port is disabled)
show platform internal counters vlan [*hw_vlan_id*]	Displays VLAN statistics.
show platform internal counters tep [*tunnel_id*]	Displays TEP statistics.
show platform internal counters flow [*rule_id* | {**dump** [*asic inst*] | [**slice direction** | **index** *hw_index*]}]	Displays flow statistics.
clear platform internal counters port [*port_num* | {**internal** [*int_port_num*]}]	Clears port statistics.
clear platform internal counters vlan [*hw_vlan_id*]	Clears VLAN counters.

Command	Purpose
debug platform internal stats logging level *log_level*	Sets the debug logging level.
debug platform internal stats logging {**err**|**trace**|**flow**}	Sets the debug logging type.

Managing Statistics Thresholds Using the GUI

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	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: • 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.	Follow these steps to resolve the issue: • 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. **Note** Except for fabric health statistics, 5 minute statistics are stored on the switch and are lost when the switch reboots.
Statistics or history are not maintained for the configured time period.	Follow these steps to resolve the issue: • 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.	Review whether the configuration exceeds the maximum historical record size. The limitations are as follows: • 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.

Problem	Solution
An export policy is configured but the APIC does not export statistics.	Follow these steps to resolve the issue: • 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 exporte 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. **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.
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.

Specifying Syslog Sources and Destinations

This section explains how to create syslog destination groups, a syslog source, and how to enable syslog to display in NX-OS CLI format using the REST API.

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.

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:

a) In the group and profile **Name** field, enter a name for the monitoring destination group and profile.
b) In the group and profile **Admin State** drop-down list, choose **enabled**.
c) 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`.
d) 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.
e) Click **Next**.
f) In the **Create Remote Destinations** area, click **+** to add a remote destination.

Caution Risk of hostname resolution failure for remote Syslog destinations, if the DNS server used 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:

a) In the **Host** field, enter an IP address or a fully qualified domain name for the destination host.
b) (Optional) In the **Name** field, enter a name for the destination host.
c) In the **Admin State** field, click the **enabled** radio button.
d) (Optional) Choose a minimum severity **Severity**, a **Port** number, and a syslog **Forwarding Facility**.
e) From the **Management EPG** drop-down list, choose the management endpoint group.
f) 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.

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:

a) Click the Edit icon to the right of the **Monitoring Object** drop-down list.
b) From the **Select Monitoring Package** drop-down list, choose an object class package.
c) Select the checkbox for each object that you want to monitor.
d) 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:

a) In the **Name** field, enter a name for the syslog source.
b) From the **Min Severity** drop-down list, choose the minimum severity of system log messages to be sent.
c) In the **Include** field, check the checkboxes for the type of messages to be sent.
d) From the **Dest Group** drop-down list, choose the syslog destination group to which the system log messages will be sent.
e) Click **Submit**.

Step 9 (Optional) To add more syslog sources, click + again and repeat the steps in the **Create Syslog Source** dialog box

Enabling Syslog to Display in NX-OS CLI Format, Using the REST API

By default the Syslog format is RFC 5424 compliant. You can change the default display of Syslogs to NX-OS type format, similar to the following example:

```
apic1# moquery -c "syslogRemoteDest"
Total Objects shown: 1

# syslog.RemoteDest
host                 : 172.23.49.77
adminState           : enabled
childAction          :
descr                :
dn                   : uni/fabric/slgroup-syslog-mpod/rdst-172.23.49.77
epgDn                :
format               : nxos
forwardingFacility   : local7
ip                   :
lcOwn                : local
modTs                : 2016-05-17T16:51:57.231-07:00
monPolDn             : uni/fabric/monfab-default
name                 : syslog-dest
operState            : unknown
port                 : 514
rn                   : rdst-172.23.49.77
severity             : information
status               :
uid                  : 15374
vrfId                : 0
vrfName              :
```

To enable the Syslogs to display in NX-OS type format, perform the following steps, using the REST API.

Step 1 Enable the Syslogs to display in NX-OS type format, as in the following example:

```
POST https://192.168.20.123/api/node/mo/uni/fabric.xml
<syslogGroup name="DestGrp77" format="nxos">
<syslogRemoteDest name="slRmtDest77" host="172.31.138.20" severity="debugging"/>
</syslogGroup>
```

The syslogGroup is the Syslog monitoring destination group, the sysLogRemoteDest is the name you previously configured for your Syslog server, and the host is the IP address for the previously configured Syslog server.

Step 2 Set the Syslog format back to the default RFC 5424 format, as in the following example:

```
POST https://192.168.20.123/api/node/mo/uni/fabric.xml
<syslogGroup name="DestGrp77" format="aci">
<syslogRemoteDest name="slRmtDest77" host="172.31.138.20" severity="debugging"/>
</syslogGroup>
```

Discovering Paths and Testing Connectivity with Traceroute

This section lists the traceroute guidelines and restriction and explains how to perform a traceroute between endpoints.

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). 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.

Performing a Traceroute Between Endpoints

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 **Troubleshoot Policies**.

Step 4 Under **Troubleshoot Policies**, right-click **Endpoint-to-Endpoint Traceroute Policies** and choose **Create Endpoint-to-Endpoint Traceroute Policy**.

Step 5 In the **Create Endpoint-to-Endpoint Traceroute Policy** dialog box, perform the following actions:

a) In the **Name** field, enter a name for the traceroute policy.
b) In the **Source End Points** table, click the + icon to edit the traceroute source.
c) From the **Source** drop-down list, choose or enter the IP address of the source endpoint and click **Update**.
d) In the **Destination End Points** table, click the + icon to edit the traceroute destination.
e) From the **Destination** drop-down list, choose or enter the IP address of the destination endpoint and click **Update**.
f) (Optional) From the **IP Protocol** drop-down list, choose a protocol for the traceroute packets. By default, the protocol is undefined (0), but it might be necessary to specify a protocol in order for the traceroute packets to traverse a filter or firewall.
g) (Optional) From the **Source Port** and **Destination Port** drop-down lists, choose source and destination protocol numbers for the traceroute packets.
h) From the **Admin State** drop-down list, choose **Start**.

i) Click **Submit** to launch the traceroute.

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 End Points** 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.

Note 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 understand and visualize how your network is behaving, which can ease your networking concerns should issues arise.

This wizard allows you (the Administrative user) to troubleshoot issues that occur during specific time frames, which can be designated by selecting two endpoints. For example, you may have two endpoints that are having intermittent packet loss but you don't understand why. Through the troubleshooting GUI, you can evaluate the issue so that you can effectively resolve it rather than logging onto each machine that you suspect to be causing this faulty behavior.

Since you may want to revisit the session later, you should give the session a unique name. You may also choose to use a pre-configured test. You can debug from endpoint to endpoint, or from an internal or external endpoint, or from an external to an internal endpoint.

Further, you can define a time window in which you want to perform the debug. The Troubleshooting GUI allows you to enter a source and destination endpoint for the endpoints you are looking for. You can do this with a MAC, IPv4, or IPv6 address and then select by tenant. You also have the option to generate a troubleshooting report that can be sent to TAC.

The following section describes the topology of the Troubleshooting Wizard, which is a simplified view of the fabric with only the elements that are relevant to your two endpoints under inspection.

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

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 Source and Destination endpoints (Eps) and select a time window for your troubleshooting session. The time window is used for retrieving Events, Fault Records, Deployment Records, Audit Logs, and Statistics. (The description and time window can only be edited in the first page of the wizard, before clicking on **Start**.)

Note You cannot modify the Source and Destination endpoints once you have clicked either the **GENERATE REPORT** button or the **START** button. If you want to change the Source and Destination information after you have entered it, you have to start a new session.

Note 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.

To set up your troubleshooting session information:

Step 1 Select **OPERATIONS** from the top of the screen, then select **VISIBILITY & TROUBLESHOOTING**, shown as follows:

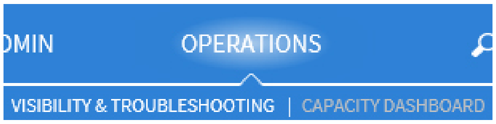

The **Visibility & Troubleshooting** screen appears.

Step 2 You can choose to either use an existing troubleshooting session (using the drop-down menu) or you can create a new one. To create a new one, enter a name for it in the **Session Name** field (shown as follows using the example, "tsw_session2").

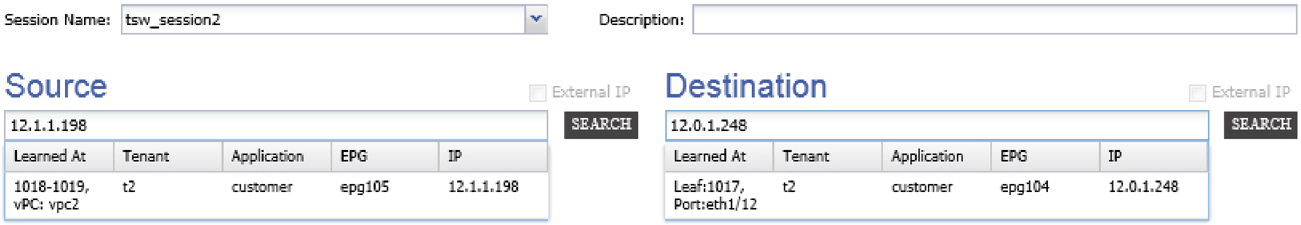

Session Name: tsw_session2

Description:

Source ☐ External IP

12.1.1.198 SEARCH

Learned At	Tenant	Application	EPG	IP
1018-1019, vPC: vpc2	t2	customer	epg105	12.1.1.198

Destination ☐ External IP

12.0.1.248 SEARCH

Learned At	Tenant	Application	EPG	IP
Leaf:1017, Port:eth1/12	t2	customer	epg104	12.0.1.248

Step 3 Enter a description in the **Description** field to provide additional information.
(This step is optional.)

Step 4 From the **Source** pull-down menu, enter a MAC, IPv4, or IPv6 address or choose an existing one.

Step 5 Click **SEARCH**.
A box appears (shown as follows) displaying one or multiple 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 number, fex number, and port details are shown in the **Learned At** column.

Step 6 From the **Destination** pull-down menu, enter a MAC, IPv4, or IPv6 address or choose an existing one.

Step 7 Click **SEARCH**.

A box appears displaying one or multiple rows to help you make a selection (as previously described for the **Source** endpoint search).

Step 8 Check the **External IP** checkbox if you are using an endpoint to external internet protocol.

Note For more information about endpoints and external IPs, refer to the *Cisco Application Centric Infrastructure Fundamentals* guide.

Note Ideally, you should select the Source and Destination endpoints from the same tenant or some of the troubleshooting functionality may be impacted, as explained later in this document. Once you make selections for these endpoints, you can learn about the topology that connects the two in the **Faults** troubleshooting screen.

Step 9 Select a time window by making selections from the **From** (for session Start time) and **To** (for session End time) pull-down menus.

The **Time Window** (shown as follows) 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 leafs (or nodes).

Note You have two options for setting the time window, which you can toggle back and forth from using the **Use fixed time** checkbox.

- You can specify a rolling time window based on any number of "Latest Minutes" (the default is 240 minutes but this can be changed), shown as follows:

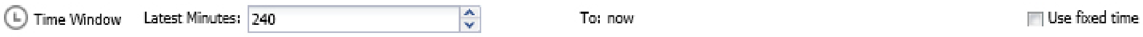

Time Window Latest Minutes: 240 To: now □ Use fixed time

- Or, you can specify a fixed time window for the session in the **From** and **To** fields by checking the **Use fixed time** checkbox, shown as follows:

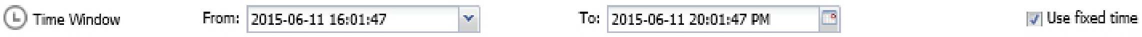

Time Window From: 2015-06-11 16:01:47 To: 2015-06-11 20:01:47 PM ☑ Use fixed time

Note The default time window is based on a default of "latest 240 minutes" (which means that the session contains data for the past 240 minutes) preceding the time you created the session. You can also set up or modify time window information from the bottom of the left navigation pane, shown as follows:

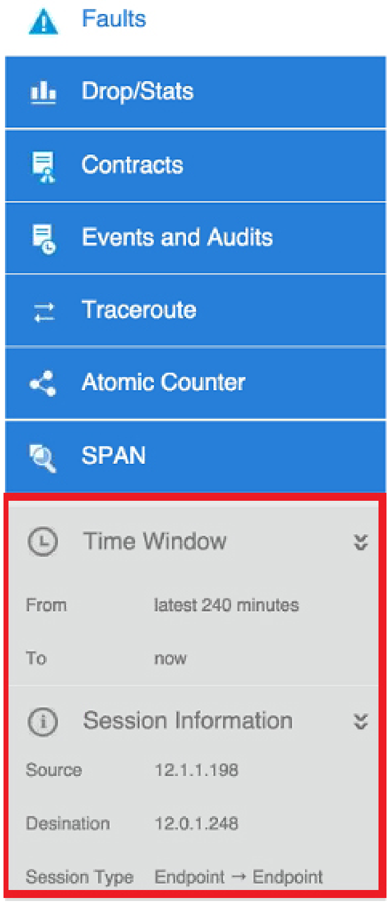

Time Window

From latest 240 minutes

To now

Session Information

Source 12.1.1.198

Desination 12.0.1.248

Session Type Endpoint → Endpoint

Step 10 Click **START** at the bottom right side of the screen to begin your troubleshooting session.
The topology diagram for your troubleshooting session loads and then appears.

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

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:

Step 1 From the bottom right corner of the screen, select **GENERATE REPORT**.
A box displays that allows you to select an output format, shown in the next step.

Step 2 Select an output format from the **Generate Report** drop-down box, (**XML, HTML, JSON,** or **PDF**), shown as follows:

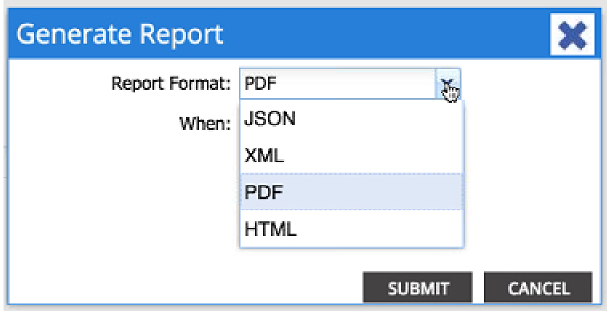

Step 3 If you want to schedule the download of the report to happen immediately, select the **Now** button, (shown as follows) then click **SUBMIT**.

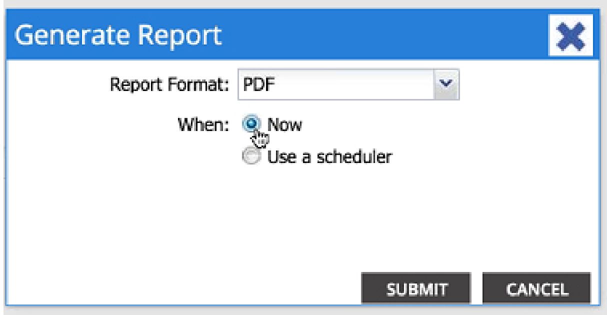

An **Information** box (shown as follows) appears indicating where to obtain the report once it has been generated.

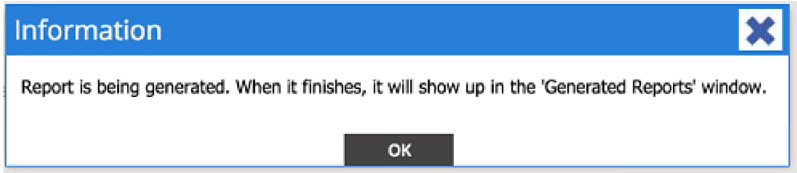

Report is being generated. When it finishes, it will show up in the 'Generated Reports' window.

Step 4 If you want to schedule the generation of the report for a later time, you can select a schedule by clicking **Use a scheduler**, shown as follows:

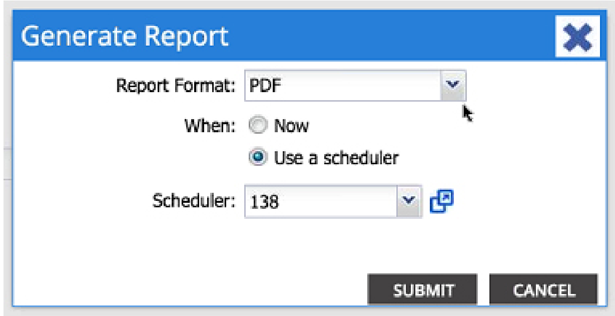

From the **Scheduler** pull-down menu, select either an existing schedule or create a new one by selecting **Create Scheduler.** shown as follows:

The **CREATE TRIGGER SCHEDULE** window appears, shown as follows:

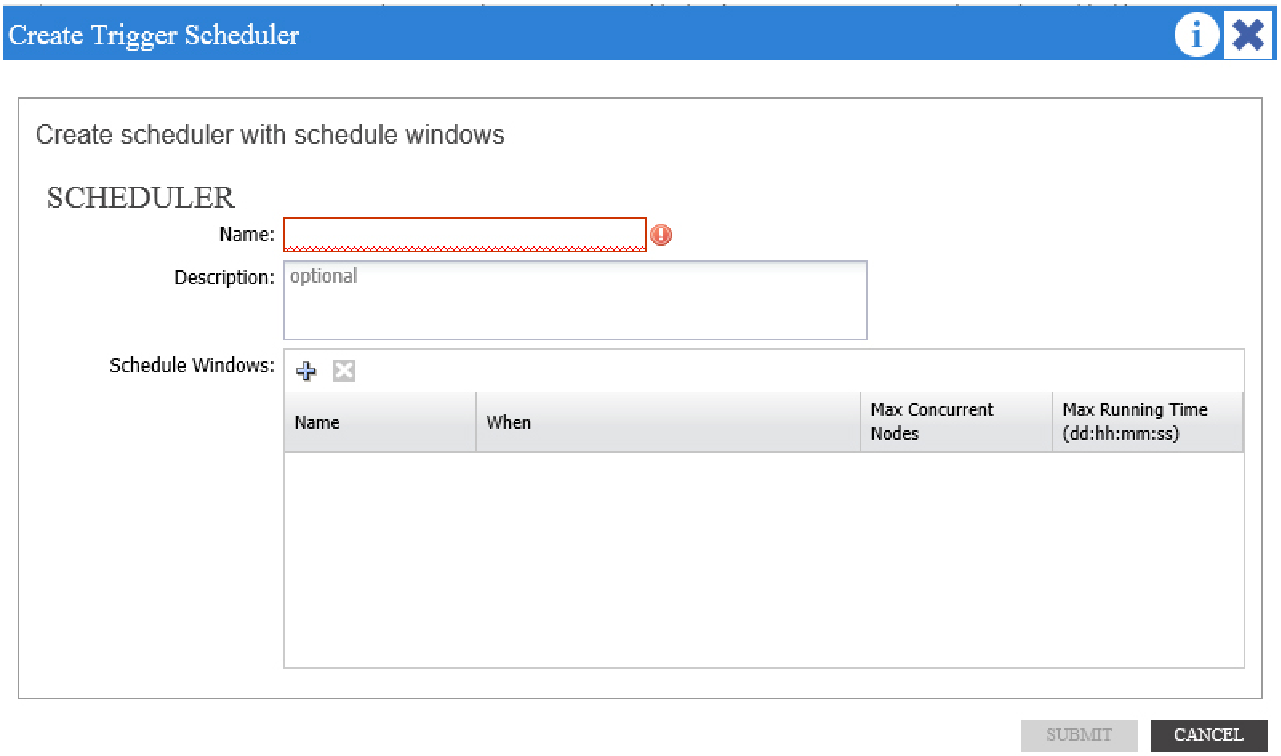

Create Trigger Scheduler

Create scheduler with schedule windows

SCHEDULER

Name:

Description: optional

Schedule Windows:

Name	When	Max Concurrent Nodes	Max Running Time (dd:hh:mm:ss)

Step 5 Enter information for **Name**, **Description** (optional), and **Schedule Windows**.

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, shown as follows:

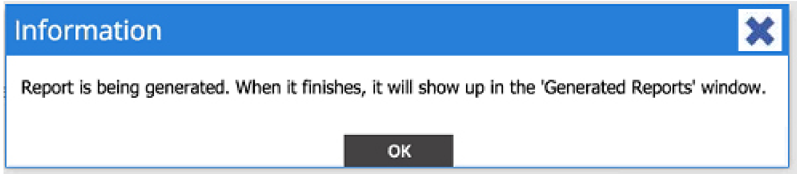

Information

Report is being generated. When it finishes, it will show up in the 'Generated Reports' window.

To retrieve and view the troubleshooting report, select **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 (shown as follows) 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 selected (for example, if the file is a PDF, it may open immediately in your browser).

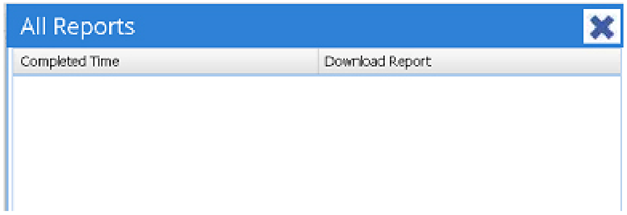

All Reports

Completed Time	Download Report

Topology in the Troubleshooting Wizard

This section explains the topology in the Troubleshooting Wizard. The topology shows how the Source and Destination end points (Eps) are connected to the fabric, what the network path is from the Source to the Destination, and what the intermediate switches are.

The Source end point is displayed on the left side of the topology and the Destination end point is on the right, as shown in the following wizard topology diagram.

Note This wizard topology only shows the leafs, spines, and fexes of the devices involved in the traffic from the Source end point to the Destination end point. However, there may be many other leafs (tens or hundreds of leafs and many other spines) that exist.

This topology also shows links, ports, and devices. If you hover over the ⓘ icon), you can see the tenant that the Ep 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).

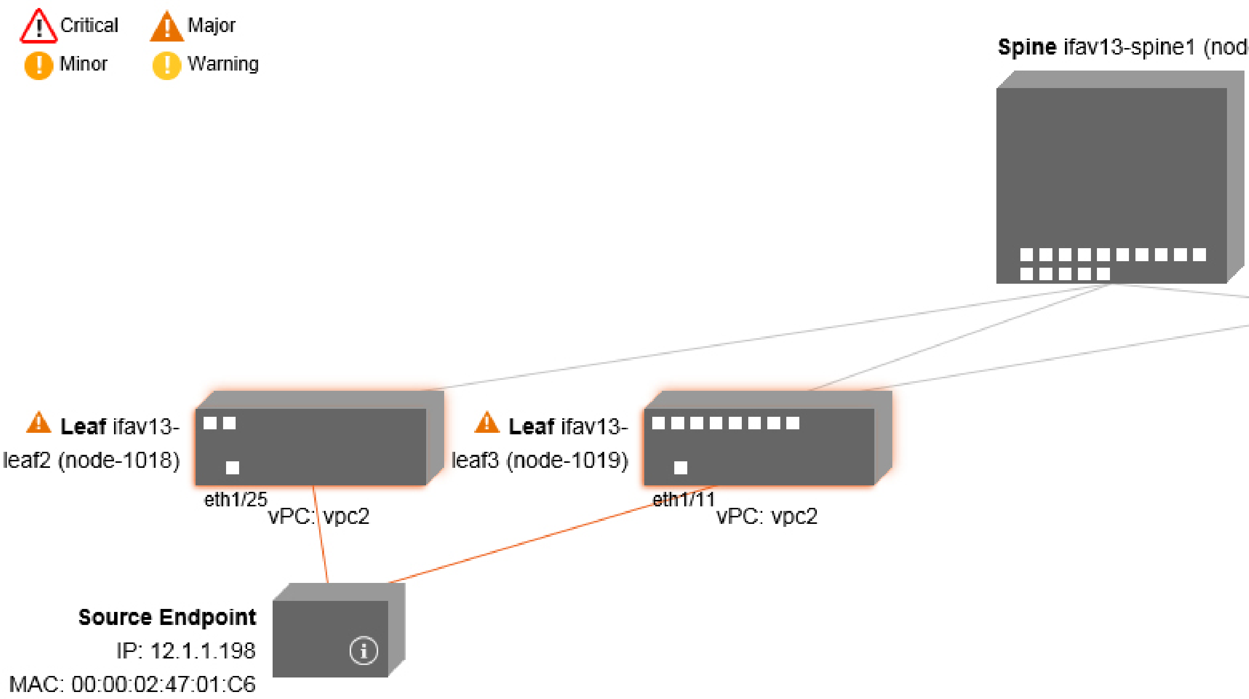

Hovering over items such as boxes or ports in the topology provides more detailed information. If the port or link has a color, this means that there is a problem for you to troubleshoot. For example, if the color is red or orange, this indicates that there is a fault on a port or link. If the color 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 Ep.

Right-clicking on a leaf 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 L4 though L7 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 VIP (Virtual IP)
- When the endpoint is behind an ESX server, the ESX is shown in the topology

Using the Faults Troubleshooting Screen

Click **Faults** in the left navigation pane (shown as follows) to begin using the **Faults** troubleshooting screen.

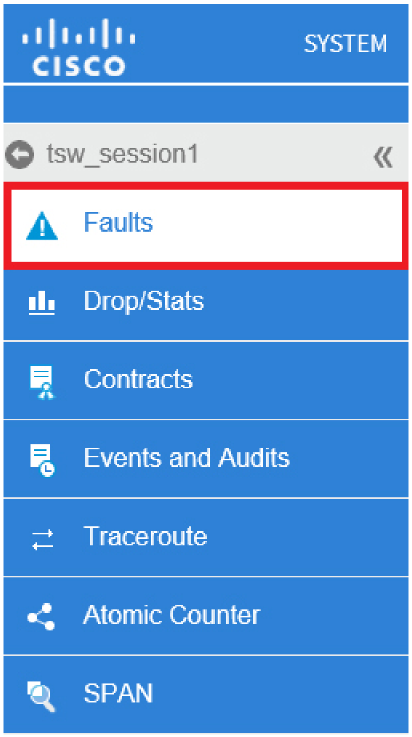

SYSTEM

tsw_session1

Faults

Drop/Stats

Contracts

Events and Audits

Traceroute

Atomic Counter

SPAN

The **Faults** screen shows the topology that connects the two endpoints 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 (shown as follows) to understand the severity levels associated with each color. This topology also shows the relevant leafs, spines, and fexes to your troubleshooting session. Hovering over items such as leafs, spines, and fexes (or clicking on faults) provides more detailed information for analysis.

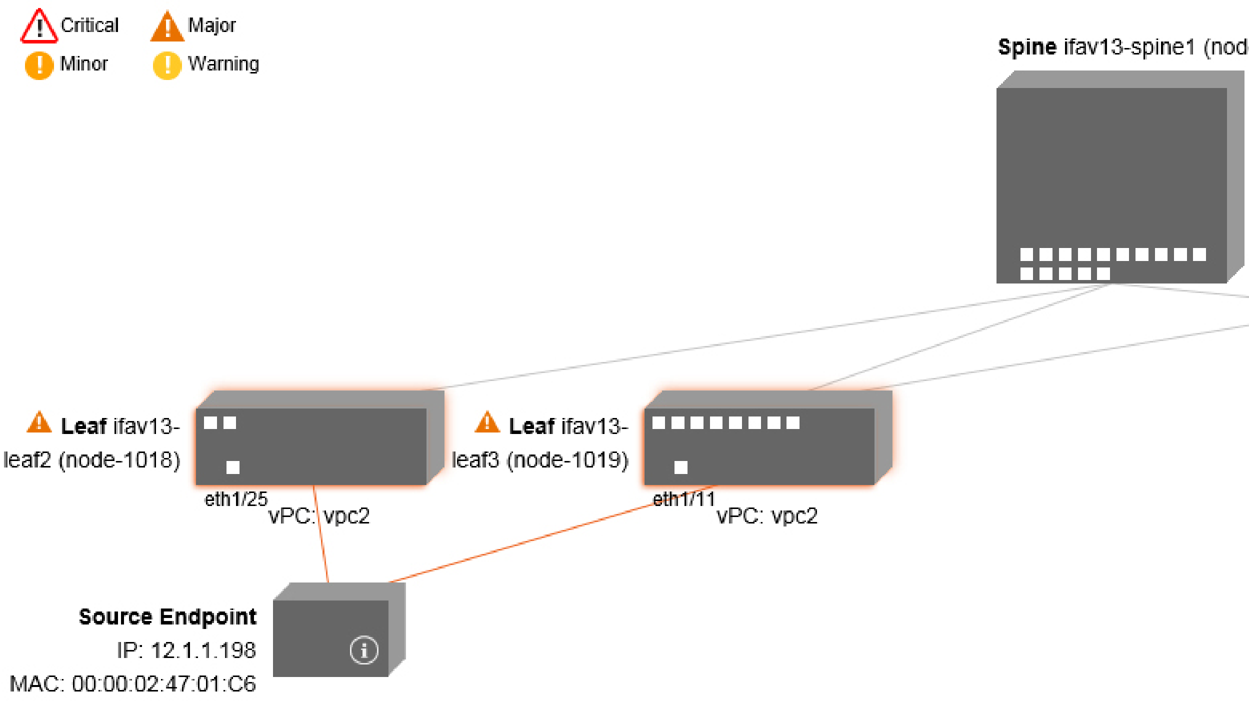

Note White boxes indicate that there are no issues to troubleshoot in that particular area.

Clicking on a fault provides a box with two tabs (**FAULTS** and **RECORDS**) that contain more detailed information for analysis, including **Severity, Affected Object, Creation Time, Last Transaction, Lifecycle**, and **Description** (shown as follows).

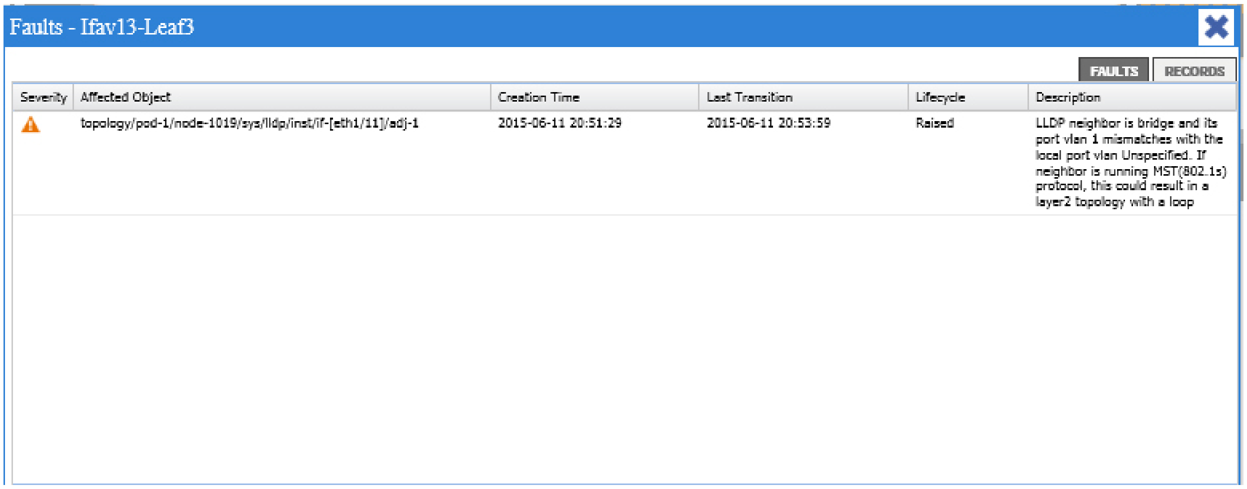

Severity	Affected Object	Creation Time	Last Transition	Lifecycle	Description
⚠	topology/pod-1/node-1019/sys/lldp/inst/if-[eth1/11]/adj-1	2015-06-11 20:51:29	2015-06-11 20:53:59	Raised	LLDP neighbor is bridge and its port vlan 1 mismatches with the local port vlan Unspecified. If neighbor is running MST(802.1s) protocol, this could result in a layer2 topology with a loop

Related Topics

Using the Drop/Statistics Troubleshooting Screen, on page 66

Using the Drop/Statistics Troubleshooting Screen

Click **Drop/Stats** in the left navigation pane (shown as follows) to begin using the **Drop/Stats** troubleshooting screen.

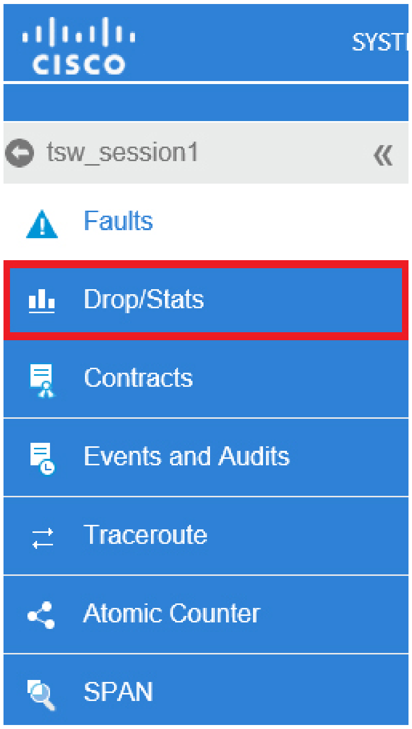

The **Drop/Stats** screen displays the topology with all the statistics from the drops so that you can clearly see where drops exist or not. You can click on any drop image to see more information for analysis.

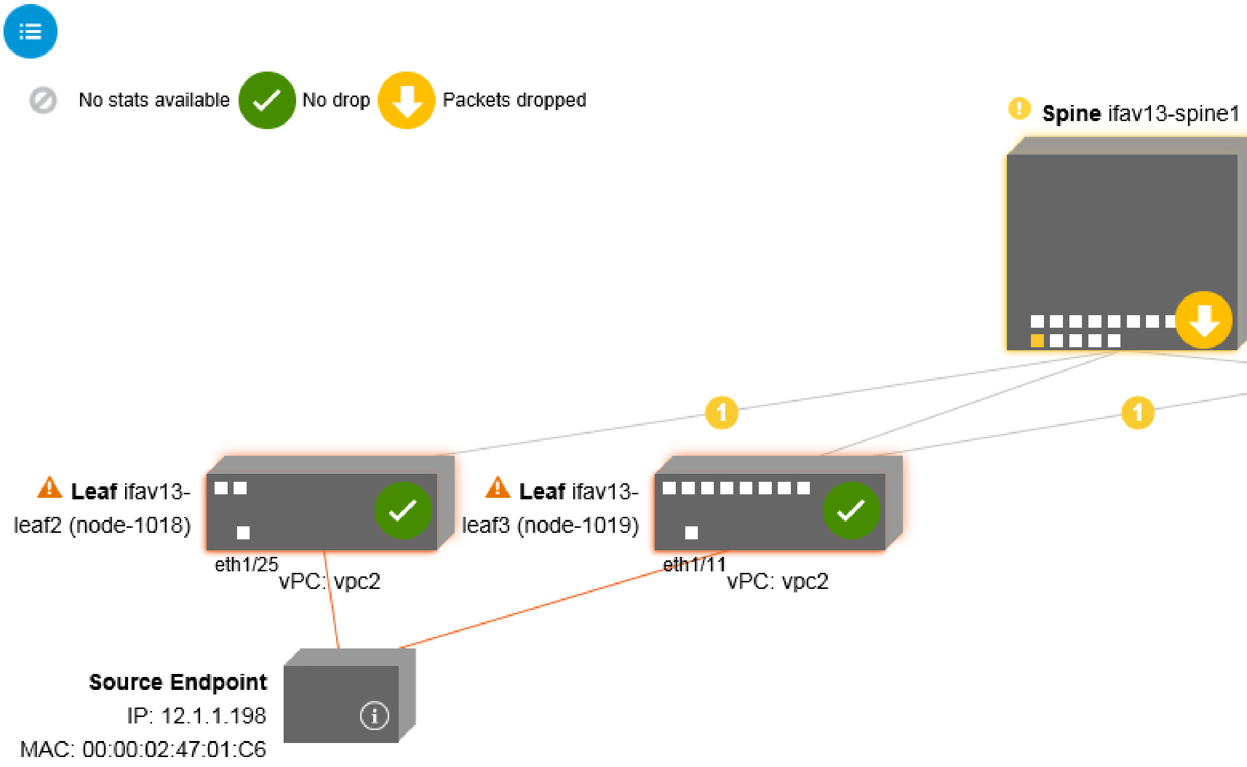

Once you click a drop image, there are three tabs at the top of the **Drop/Stats** screen, and the statistics shown are localized to that particular leaf or switch.

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, shown as follows:

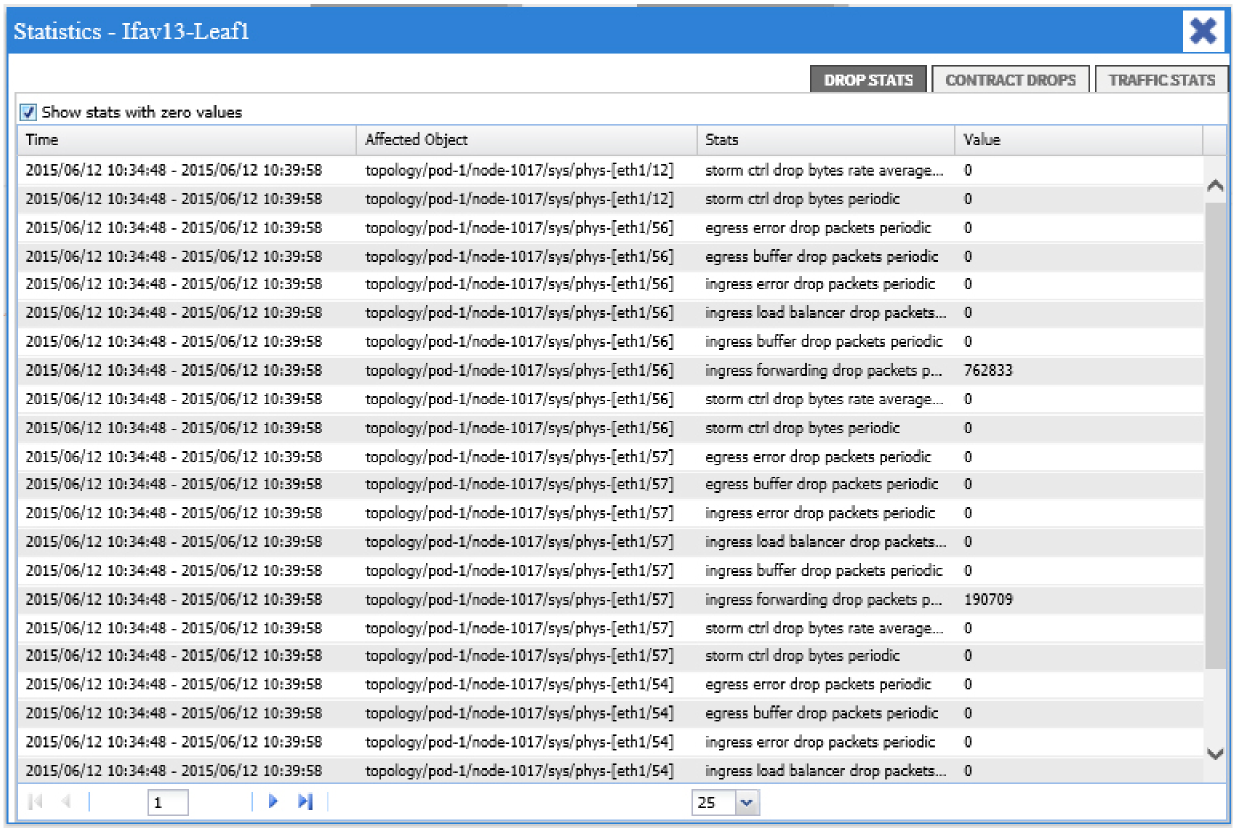

Statistics - Ifav13-Leaf1

☑ Show stats with zero values

Time	Affected Object	Stats	Value
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/12]	storm ctrl drop bytes rate average...	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/12]	storm ctrl drop bytes periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/56]	egress error drop packets periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/56]	egress buffer drop packets periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/56]	ingress error drop packets periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/56]	ingress load balancer drop packets...	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/56]	ingress buffer drop packets periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/56]	ingress forwarding drop packets p...	762833
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/56]	storm ctrl drop bytes rate average...	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/56]	storm ctrl drop bytes periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/57]	egress error drop packets periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/57]	egress buffer drop packets periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/57]	ingress error drop packets periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/57]	ingress load balancer drop packets...	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/57]	ingress buffer drop packets periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/57]	ingress forwarding drop packets p...	190709
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/57]	storm ctrl drop bytes rate average...	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/57]	storm ctrl drop bytes periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/54]	egress error drop packets periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/54]	egress buffer drop packets periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/54]	ingress error drop packets periodic	0
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/54]	ingress load balancer drop packets...	0

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

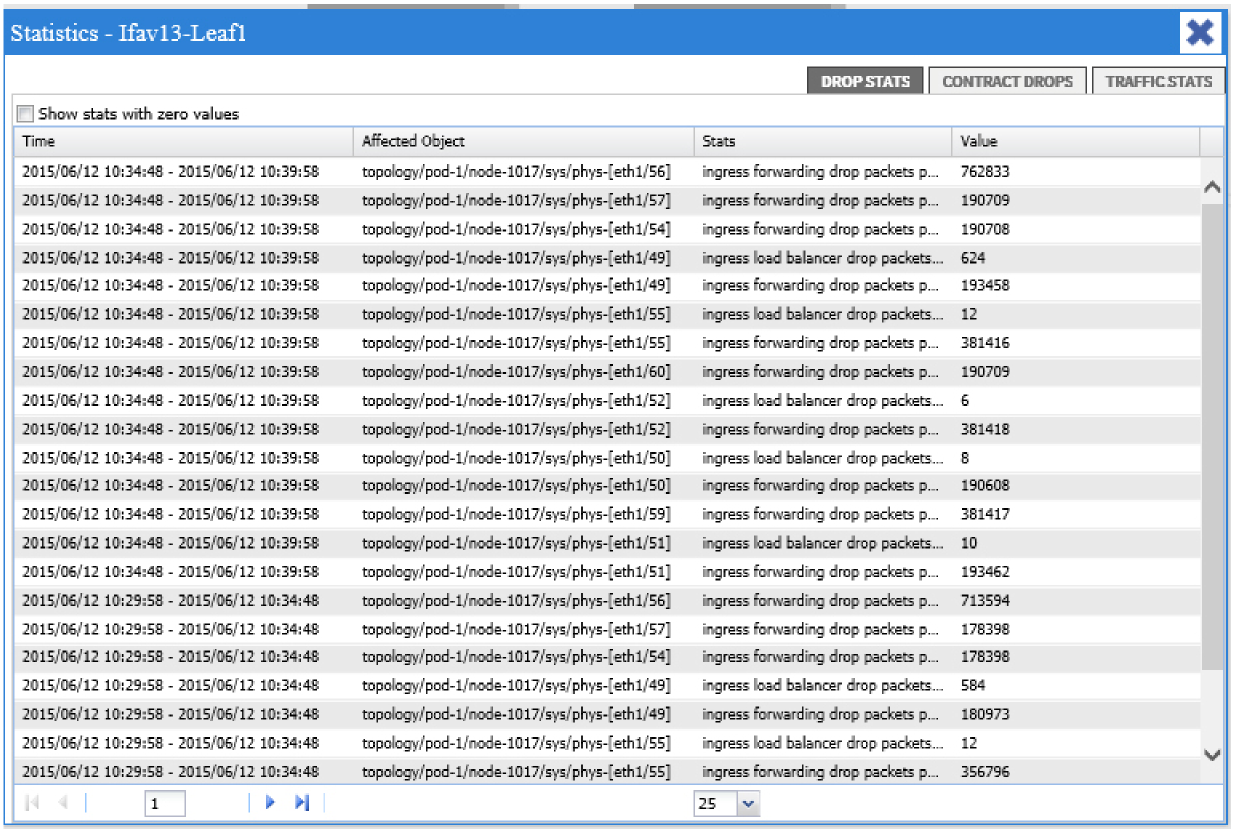

Statistics - Ifav13-Leaf1

☐ Show stats with zero values

Time	Affected Object	Stats	Value
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/56]	ingress forwarding drop packets p...	762833
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/57]	ingress forwarding drop packets p...	190709
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/54]	ingress forwarding drop packets p...	190708
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/49]	ingress load balancer drop packets...	624
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/49]	ingress forwarding drop packets p...	193458
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/55]	ingress load balancer drop packets...	12
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/55]	ingress forwarding drop packets p...	381416
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/60]	ingress forwarding drop packets p...	190709
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/52]	ingress load balancer drop packets...	6
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/52]	ingress forwarding drop packets p...	381418
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/50]	ingress load balancer drop packets...	8
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/50]	ingress forwarding drop packets p...	190608
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/59]	ingress forwarding drop packets p...	381417
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/51]	ingress load balancer drop packets...	10
2015/06/12 10:34:48 - 2015/06/12 10:39:58	topology/pod-1/node-1017/sys/phys-[eth1/51]	ingress forwarding drop packets p...	193462
2015/06/12 10:29:58 - 2015/06/12 10:34:48	topology/pod-1/node-1017/sys/phys-[eth1/56]	ingress forwarding drop packets p...	713594
2015/06/12 10:29:58 - 2015/06/12 10:34:48	topology/pod-1/node-1017/sys/phys-[eth1/57]	ingress forwarding drop packets p...	178398
2015/06/12 10:29:58 - 2015/06/12 10:34:48	topology/pod-1/node-1017/sys/phys-[eth1/54]	ingress forwarding drop packets p...	178398
2015/06/12 10:29:58 - 2015/06/12 10:34:48	topology/pod-1/node-1017/sys/phys-[eth1/49]	ingress load balancer drop packets...	584
2015/06/12 10:29:58 - 2015/06/12 10:34:48	topology/pod-1/node-1017/sys/phys-[eth1/49]	ingress forwarding drop packets p...	180973
2015/06/12 10:29:58 - 2015/06/12 10:34:48	topology/pod-1/node-1017/sys/phys-[eth1/55]	ingress load balancer drop packets...	12
2015/06/12 10:29:58 - 2015/06/12 10:34:48	topology/pod-1/node-1017/sys/phys-[eth1/55]	ingress forwarding drop packets p...	356796

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. The following shows the **Statistics - All** screen with **Show stats with zero values** checked.

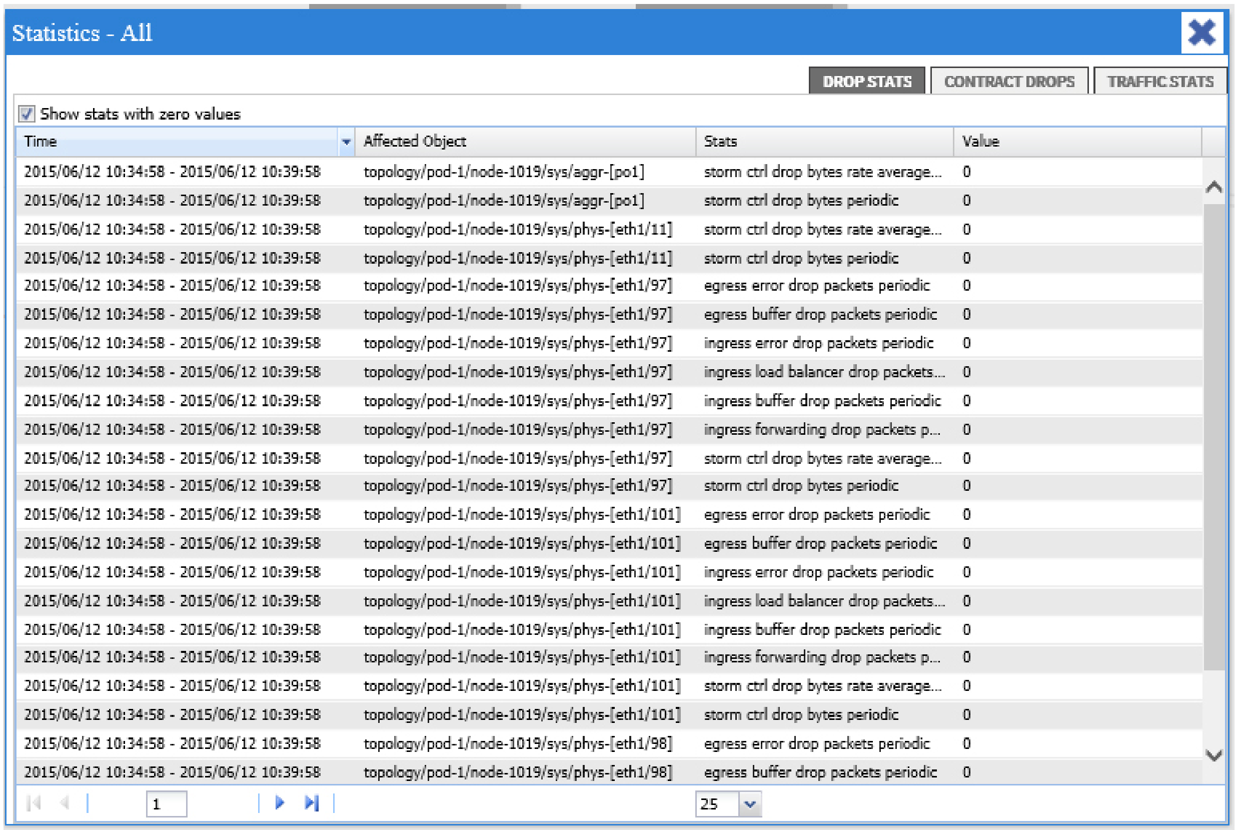

☑ Show stats with zero values

Time	Affected Object	Stats	Value
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/aggr-[po1]	storm ctrl drop bytes rate average...	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/aggr-[po1]	storm ctrl drop bytes periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/11]	storm ctrl drop bytes rate average...	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/11]	storm ctrl drop bytes periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/97]	egress error drop packets periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/97]	egress buffer drop packets periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/97]	ingress error drop packets periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/97]	ingress load balancer drop packets...	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/97]	ingress buffer drop packets periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/97]	ingress forwarding drop packets p...	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/97]	storm ctrl drop bytes rate average...	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/97]	storm ctrl drop bytes periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/101]	egress error drop packets periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/101]	egress buffer drop packets periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/101]	ingress error drop packets periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/101]	ingress load balancer drop packets...	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/101]	ingress buffer drop packets periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/101]	ingress forwarding drop packets p...	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/101]	storm ctrl drop bytes rate average...	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/101]	storm ctrl drop bytes periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/98]	egress error drop packets periodic	0
2015/06/12 10:34:58 - 2015/06/12 10:39:58	topology/pod-1/node-1019/sys/phys-[eth1/98]	egress buffer drop packets periodic	0

Related Topics

Using the Contracts Troubleshooting Screen, on page 70

Using the Contracts Troubleshooting Screen

Click **Contracts** in the left navigation pane (shown as follows) to begin using the **Contracts** troubleshooting screen.

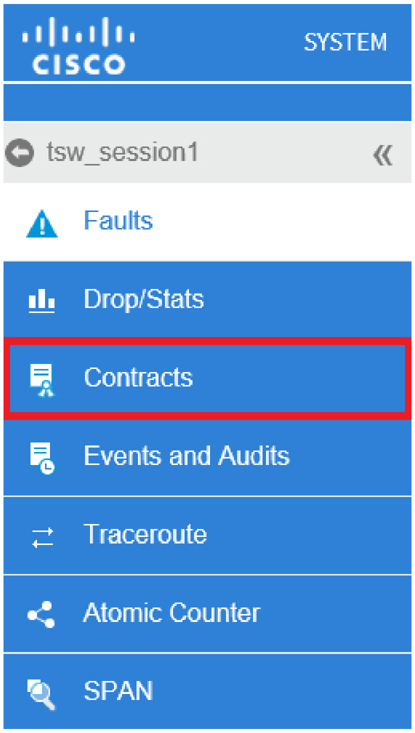

The **Contracts** troubleshooting screen displays the contracts that are applicable from the Source to the Destination and from the Destination to the Source, shown as follows:

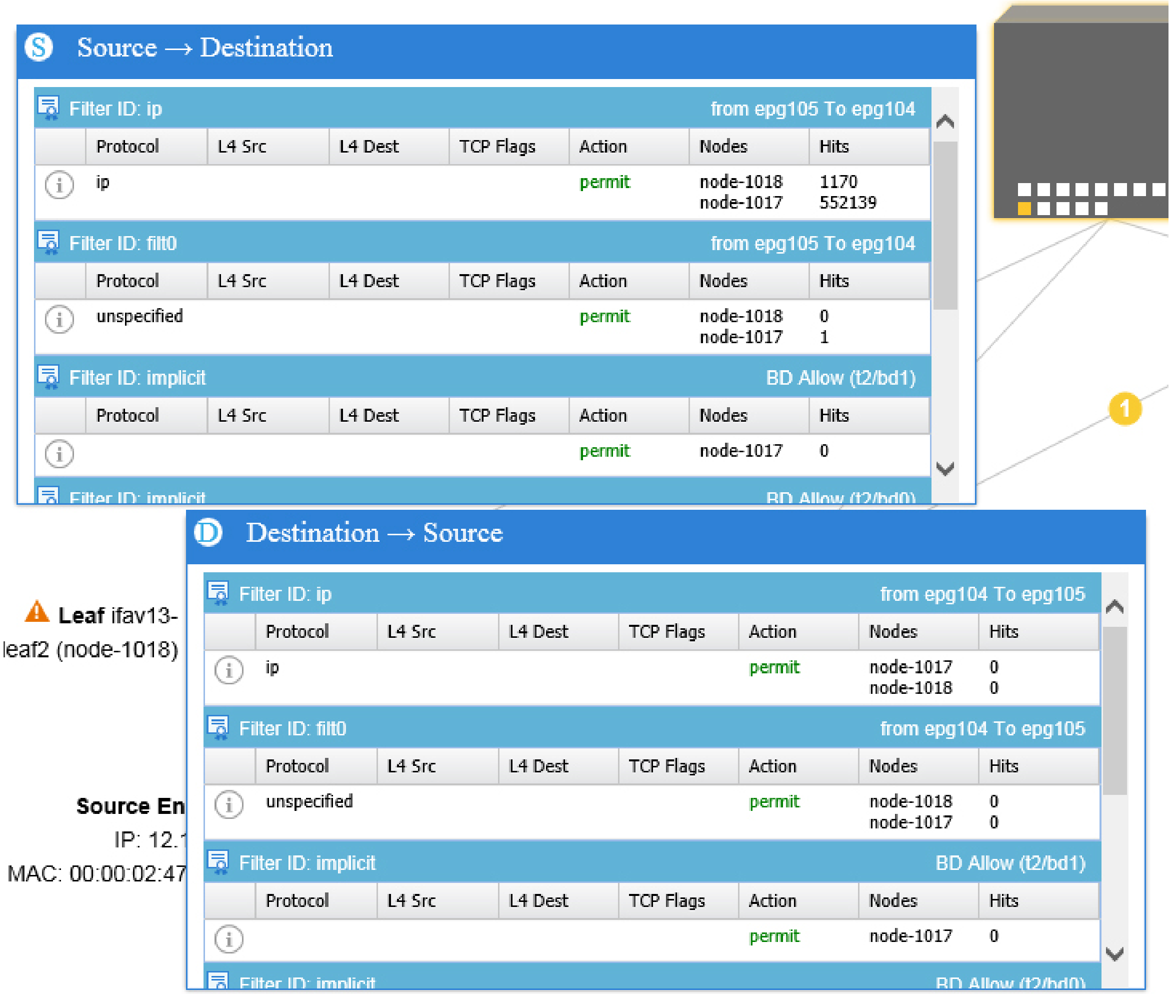

Source → Destination

Filter ID: ip — from epg105 To epg104

	Protocol	L4 Src	L4 Dest	TCP Flags	Action	Nodes	Hits
ⓘ	ip				permit	node-1018 node-1017	1170 552139

Filter ID: filt0 — from epg105 To epg104

	Protocol	L4 Src	L4 Dest	TCP Flags	Action	Nodes	Hits
ⓘ	unspecified				permit	node-1018 node-1017	0 1

Filter ID: implicit — BD Allow (t2/bd1)

	Protocol	L4 Src	L4 Dest	TCP Flags	Action	Nodes	Hits
ⓘ					permit	node-1017	0

Filter ID: implicit — BD Allow (t2/bd0)

Destination → Source

Filter ID: ip — from epg104 To epg105

	Protocol	L4 Src	L4 Dest	TCP Flags	Action	Nodes	Hits
ⓘ	ip				permit	node-1017 node-1018	0 0

Filter ID: filt0 — from epg104 To epg105

	Protocol	L4 Src	L4 Dest	TCP Flags	Action	Nodes	Hits
ⓘ	unspecified				permit	node-1018 node-1017	0 0

Filter ID: implicit — BD Allow (t2/bd1)

	Protocol	L4 Src	L4 Dest	TCP Flags	Action	Nodes	Hits
ⓘ					permit	node-1017	0

Filter ID: implicit — BD Allow (t2/bd0)

Each one of the blue table heading rows (shown above) 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 (shown as follows), shows you the contract name and the contract filter name.

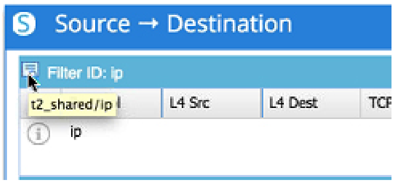

Source → Destination

Filter ID: ip

t2_shared/ip	L4 Src	L4 Dest	TCF
ⓘ ip			

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 73

Using the Events Troubleshooting Screen

Click **Events and Audits** in the left navigation pane (shown as follows) to begin using the **Events and Audits** troubleshooting screen.

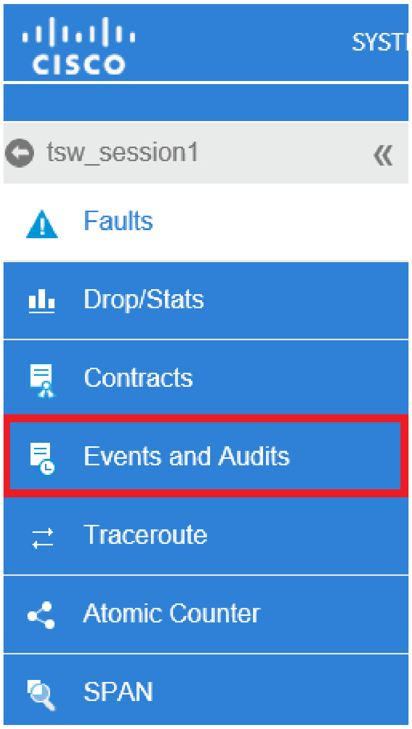

If you click on an individual leaf (shown as follows) or spine, you can see more detailed information about that individual event.

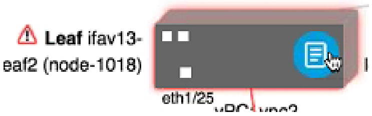

An example of this individual event information is shown as follows:

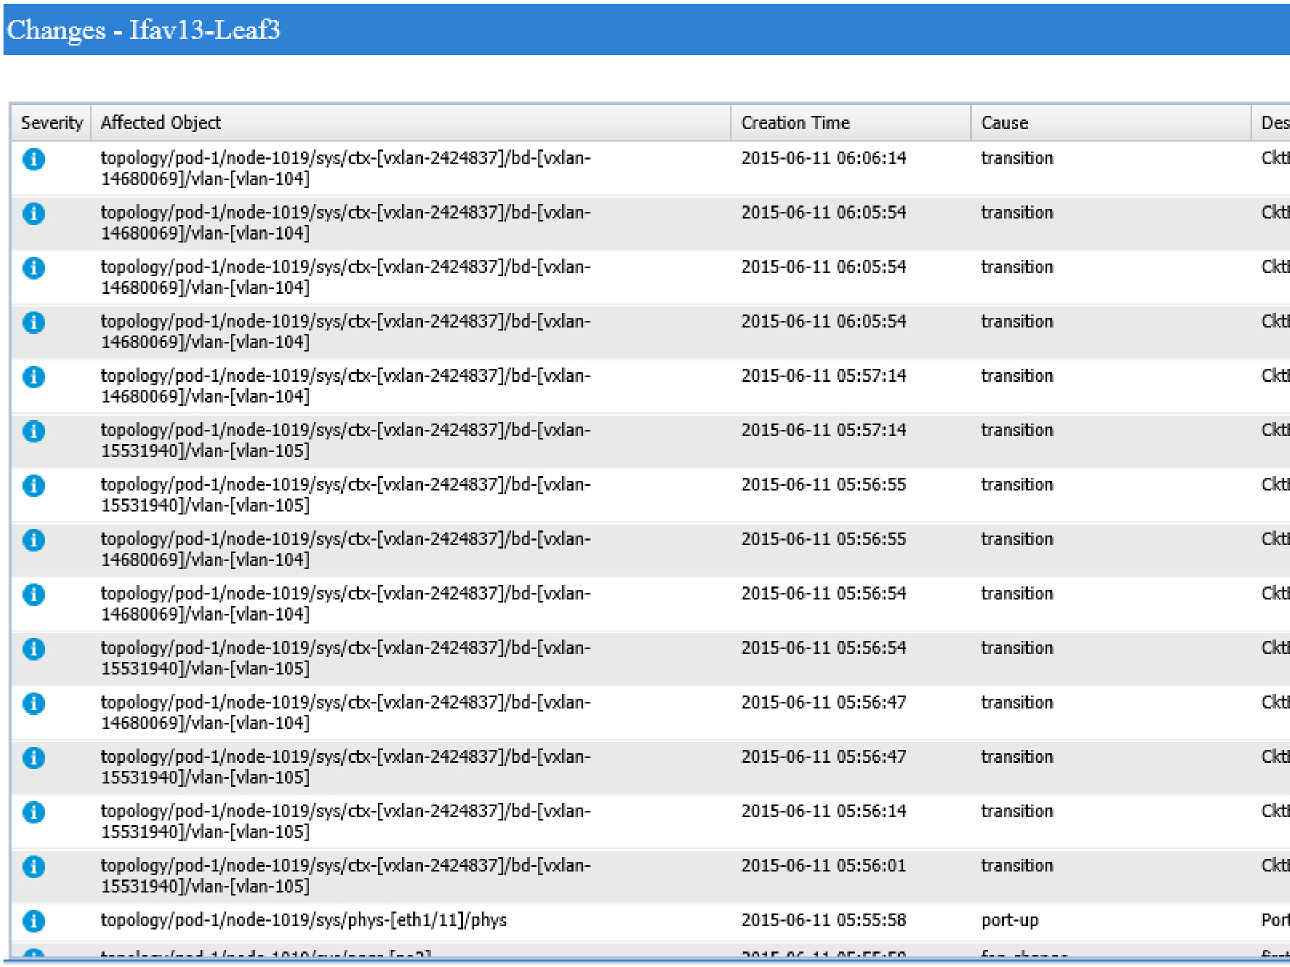

Changes - Ifav13-Leaf3

Severity	Affected Object	Creation Time	Cause	Des
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-14680069]/vlan-[vlan-104]	2015-06-11 06:06:14	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-14680069]/vlan-[vlan-104]	2015-06-11 06:05:54	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-14680069]/vlan-[vlan-104]	2015-06-11 06:05:54	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-14680069]/vlan-[vlan-104]	2015-06-11 06:05:54	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-14680069]/vlan-[vlan-104]	2015-06-11 05:57:14	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-15531940]/vlan-[vlan-105]	2015-06-11 05:57:14	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-15531940]/vlan-[vlan-105]	2015-06-11 05:56:55	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-14680069]/vlan-[vlan-104]	2015-06-11 05:56:55	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-14680069]/vlan-[vlan-104]	2015-06-11 05:56:54	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-15531940]/vlan-[vlan-105]	2015-06-11 05:56:54	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-14680069]/vlan-[vlan-104]	2015-06-11 05:56:47	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-15531940]/vlan-[vlan-105]	2015-06-11 05:56:47	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-15531940]/vlan-[vlan-105]	2015-06-11 05:56:14	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/ctx-[vxlan-2424837]/bd-[vxlan-15531940]/vlan-[vlan-105]	2015-06-11 05:56:01	transition	Ckt
ℹ	topology/pod-1/node-1019/sys/phys-[eth1/11]/phys	2015-06-11 05:55:58	port-up	Port

The two tabs available on this screen are **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 select 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):

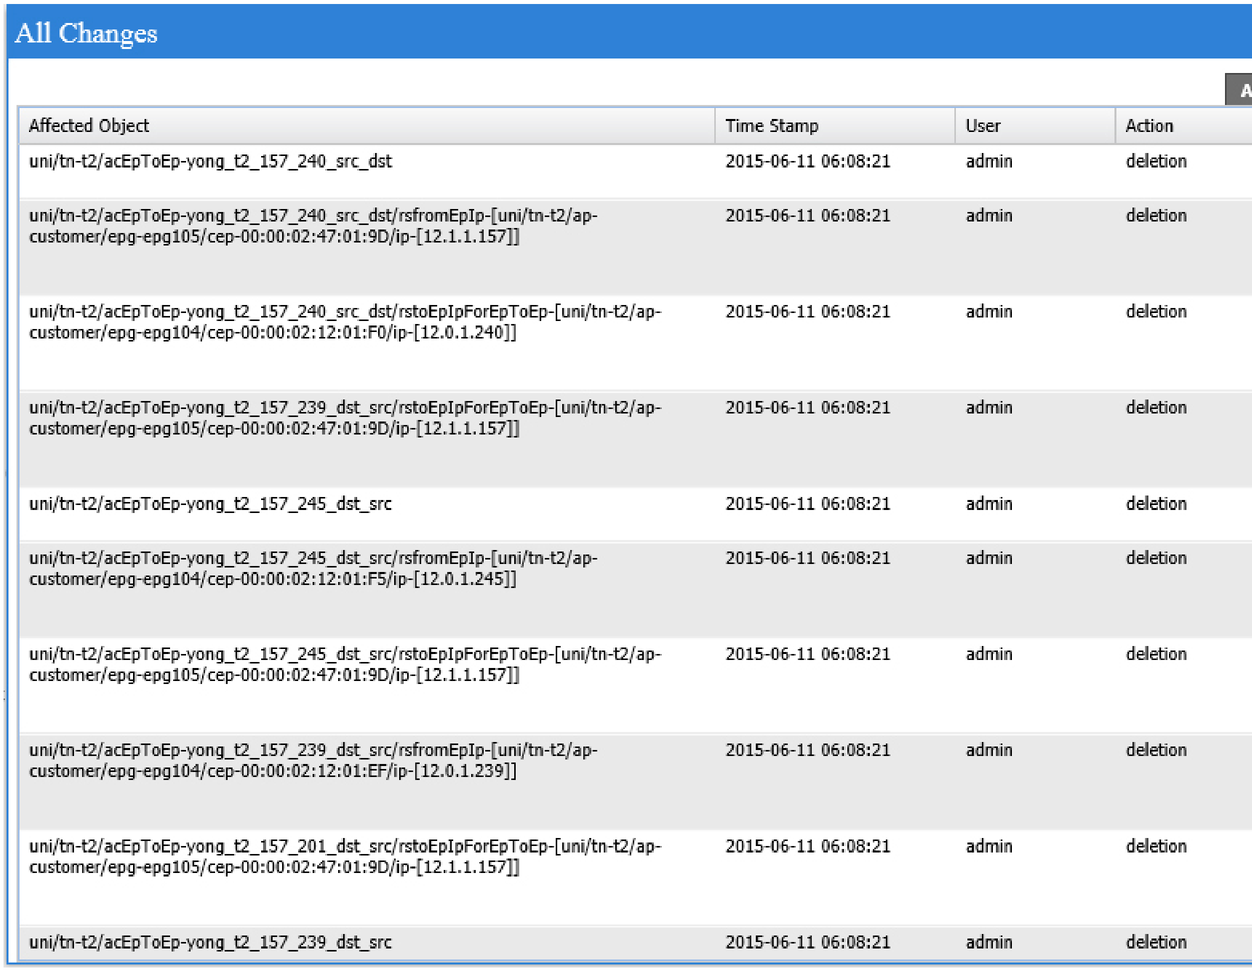

Affected Object	Time Stamp	User	Action
uni/tn-t2/acEpToEp-yong_t2_157_240_src_dst	2015-06-11 06:08:21	admin	deletion
uni/tn-t2/acEpToEp-yong_t2_157_240_src_dst/rsfromEpIp-[uni/tn-t2/ap-customer/epg-epg105/cep-00:00:02:47:01:9D/ip-[12.1.1.157]]	2015-06-11 06:08:21	admin	deletion
uni/tn-t2/acEpToEp-yong_t2_157_240_src_dst/rstoEpIpForEpToEp-[uni/tn-t2/ap-customer/epg-epg104/cep-00:00:02:12:01:F0/ip-[12.0.1.240]]	2015-06-11 06:08:21	admin	deletion
uni/tn-t2/acEpToEp-yong_t2_157_239_dst_src/rstoEpIpForEpToEp-[uni/tn-t2/ap-customer/epg-epg105/cep-00:00:02:47:01:9D/ip-[12.1.1.157]]	2015-06-11 06:08:21	admin	deletion
uni/tn-t2/acEpToEp-yong_t2_157_245_dst_src	2015-06-11 06:08:21	admin	deletion
uni/tn-t2/acEpToEp-yong_t2_157_245_dst_src/rsfromEpIp-[uni/tn-t2/ap-customer/epg-epg104/cep-00:00:02:12:01:F5/ip-[12.0.1.245]]	2015-06-11 06:08:21	admin	deletion
uni/tn-t2/acEpToEp-yong_t2_157_245_dst_src/rstoEpIpForEpToEp-[uni/tn-t2/ap-customer/epg-epg105/cep-00:00:02:47:01:9D/ip-[12.1.1.157]]	2015-06-11 06:08:21	admin	deletion
uni/tn-t2/acEpToEp-yong_t2_157_239_dst_src/rsfromEpIp-[uni/tn-t2/ap-customer/epg-epg104/cep-00:00:02:12:01:EF/ip-[12.0.1.239]]	2015-06-11 06:08:21	admin	deletion
uni/tn-t2/acEpToEp-yong_t2_157_201_dst_src/rstoEpIpForEpToEp-[uni/tn-t2/ap-customer/epg-epg105/cep-00:00:02:47:01:9D/ip-[12.1.1.157]]	2015-06-11 06:08:21	admin	deletion
uni/tn-t2/acEpToEp-yong_t2_157_239_dst_src	2015-06-11 06:08:21	admin	deletion

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 76

Using the Traceroute Troubleshooting Screen

Click **Traceroute** in the left navigation pane (shown as follows) to begin using the **Traceroute** troubleshooting screen.

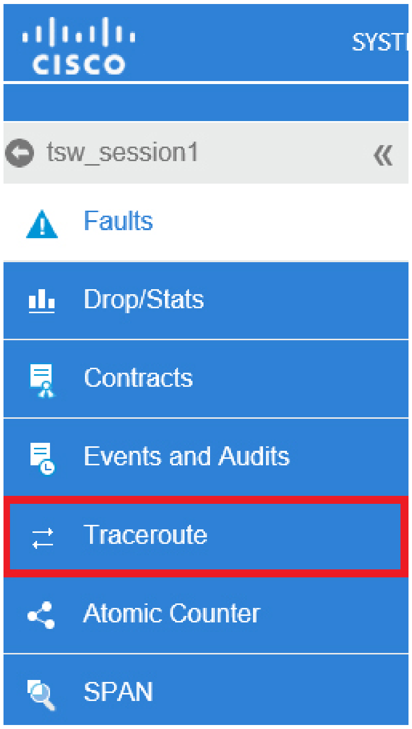

To create and run a traceroute for troubleshooting:

1 In the **TRACEROUTE** box, select a destination port from the **Destination Port** pull-down menu.

2 Select a protocol from the **Protocol** pull-down menu, shown as follows:

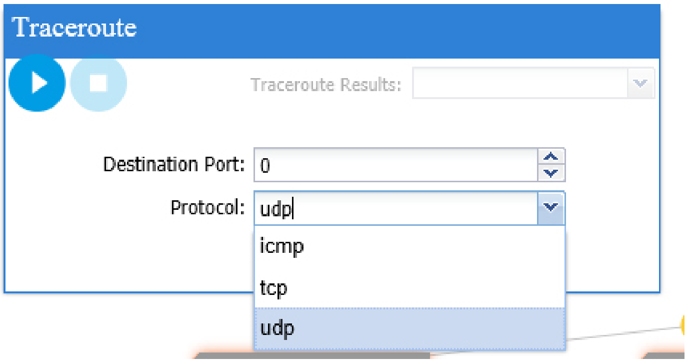

Destination Port: 0

Protocol: udp

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, shown as follows:

Note When you press the **Play** button, the polices are created on the system.

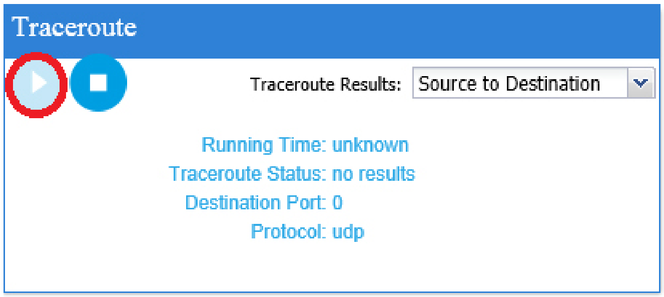

Traceroute Results: Source to Destination

Running Time: unknown
Traceroute Status: no results
Destination Port: 0
Protocol: udp

Note A **Warning** message appears, shown as follows:

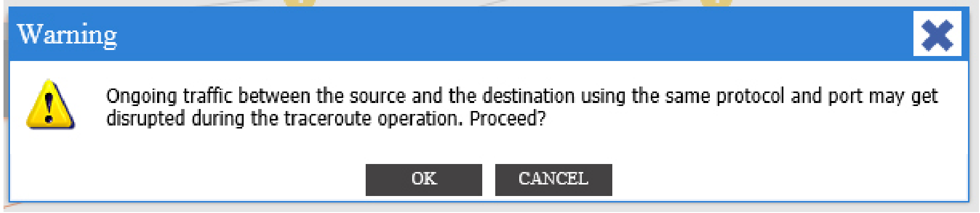

Ongoing traffic between the source and the destination using the same protocol and port may get disrupted during the traceroute operation. Proceed?

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), shown as follows:

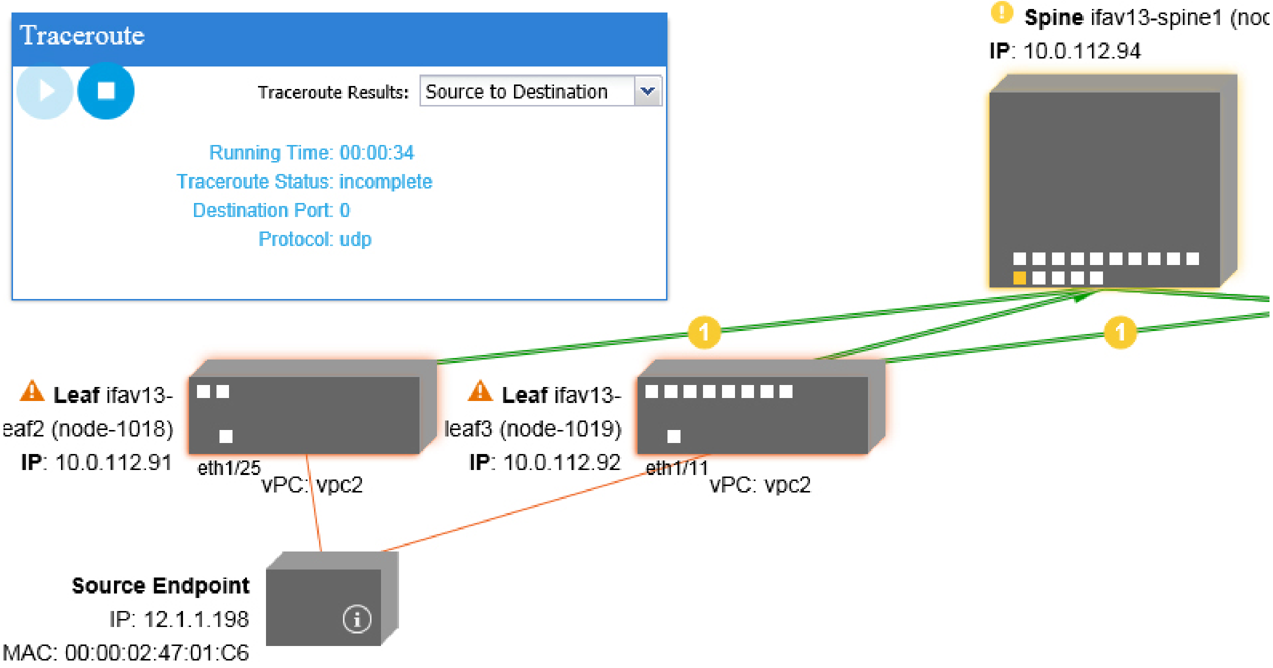

Traceroute

Traceroute Results: Source to Destination

Running Time: 00:00:34

Traceroute Status: incomplete

Destination Port: 0

Protocol: udp

The result is also shown in the **Traceroute** box (shown above), 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.

Note 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 79

Using the Atomic Counter Troubleshooting Screen

Click **Atomic Counter** in the left navigation pane (shown as follows) to begin using the **Atomic Counter** troubleshooting screen.

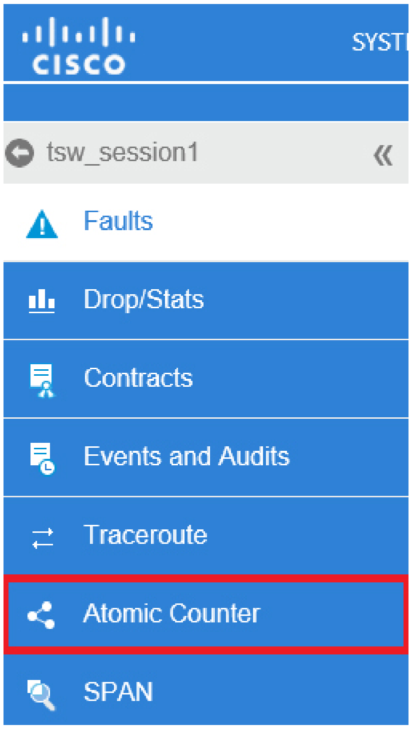

The Atomic Counter screen is used to take source and destination information and create a counter policy based on that. You can create an atomic counter policy between two endpoints and monitor the traffic going back and forth from the Source to the Destination and from the Destination to the Source. 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 (shown as follows)

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

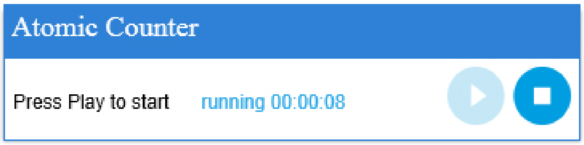

Atomic Counter

Press Play to start running 00:00:08

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.

The brief format is shown as follows:

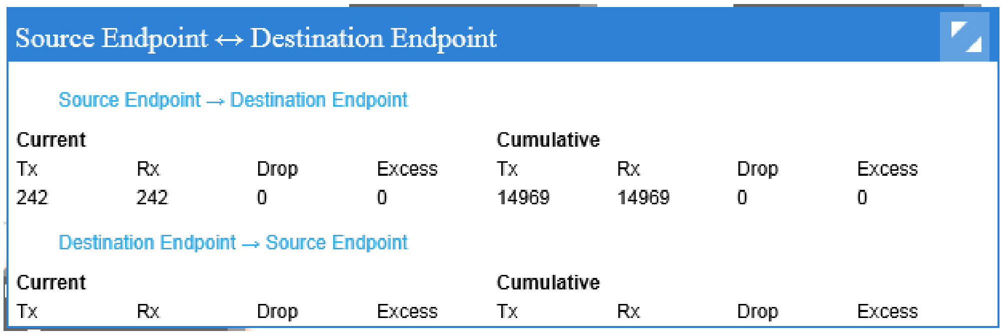

Source Endpoint ↔ Destination Endpoint

Source Endpoint → Destination Endpoint

Current				Cumulative			
Tx	Rx	Drop	Excess	Tx	Rx	Drop	Excess
242	242	0	0	14969	14969	0	0

Destination Endpoint → Source Endpoint

Current				Cumulative			
Tx	Rx	Drop	Excess	Tx	Rx	Drop	Excess

The expanded format is shown as follows:

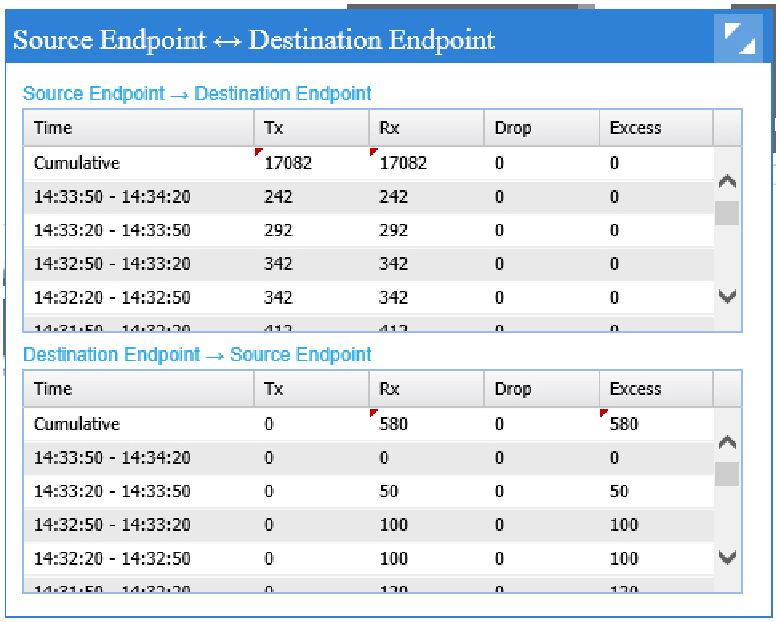

Source Endpoint ↔ Destination Endpoint

Source Endpoint → Destination Endpoint

Time	Tx	Rx	Drop	Excess
Cumulative	17082	17082	0	0
14:33:50 - 14:34:20	242	242	0	0
14:33:20 - 14:33:50	292	292	0	0
14:32:50 - 14:33:20	342	342	0	0
14:32:20 - 14:32:50	342	342	0	0
14:31:50 - 14:32:20	412	412	0	0

Destination Endpoint → Source Endpoint

Time	Tx	Rx	Drop	Excess
Cumulative	0	580	0	580
14:33:50 - 14:34:20	0	0	0	0
14:33:20 - 14:33:50	0	50	0	50
14:32:50 - 14:33:20	0	100	0	100
14:32:20 - 14:32:50	0	100	0	100
14:31:50 - 14:32:20	0	120	0	120

Related Topics

Using the SPAN Troubleshooting Screen, on page 81

Using the SPAN Troubleshooting Screen

Click **SPAN** in the left navigation pane (shown as follows) to begin using the **SPAN** troubleshooting screen.

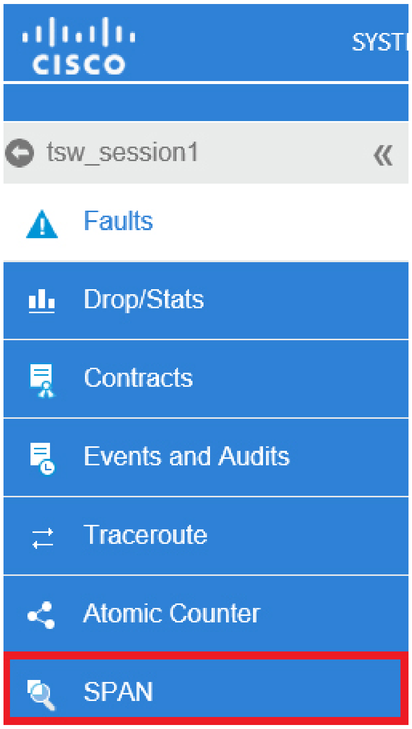

Using this screen, you can span (or mirror) bi-directional traffic and redirect it to the analyzer. In a SPAN session, you are making a copy and sending it to the analyzer.

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.

The SPAN - Bidirectional ERSPAN box is shown as follows:

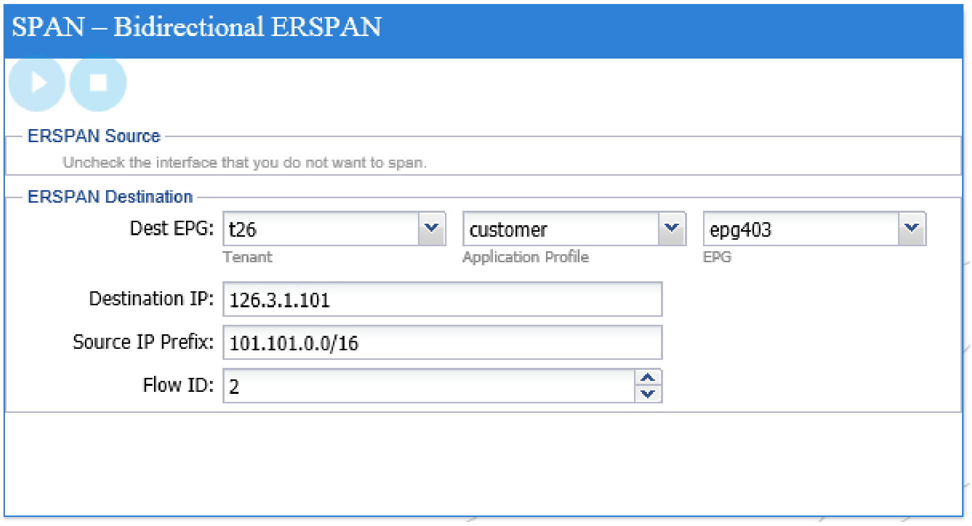

SPAN – Bidirectional ERSPAN

ERSPAN Source

Uncheck the interface that you do not want to span.

ERSPAN Destination

Dest EPG: t26 (Tenant) customer (Application Profile) epg403 (EPG)

Destination IP: 126.3.1.101

Source IP Prefix: 101.101.0.0/16

Flow ID: 2

Note When you press the **Play** button, the polices 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*.

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)

Scenario	1	2	3	4	5	6
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 84
- createsession API, on page 85
- modifysession API, on page 86
- atomiccounter API, on page 86
- traceroute API, on page 87
- span API, on page 87
- generatereport API, on page 89
- schedulereport API, on page 89
- getreportstatus API, on page 90
- getreportslist API, on page 90
- getsessionslist API, on page 90
- getsessiondetail API, on page 90
- deletesession API, on page 91
- clearreports API, on page 92
- contracts API, on page 92

interactive API

To create an endpoint (ep) to endpoint interactive troubleshooting session, use the **interactive** API. The module name is **troubleshoot.eptoeputils.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 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
- 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.eptoeputils.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.eptoeputils.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
- 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
- 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.eptoeputils.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.eptoeputils.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 93
- modifysession API, on page 94
- atomiccounter API, on page 95
- traceroute API, on page 96
- span API, on page 97
- generatereport API, on page 98
- schedulereport API, on page 99
- getreportstatus API, on page 90
- getreportslist API, on page 90
- clearreports API, on page 92
- createsession API, on page 93
- getsessionslist API, on page 101

- getsessiondetail API, on page 102
- deletesession API, on page 103
- contracts API, on page 104
- ratelimit API, on page 104
- 13ext API, on page 105

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
- 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

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 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
- dstIp	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
- portslist	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
- spansrcports	Span source ports

generatereport API

To generate a troubleshooting report using the API, use the **generatereport** API. The module name is **troubleshoot.eptoeputils.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

schedulereport API

To schedule the generation of a troubleshooting report using the API, use the **schedulereport** API. The module name is **troubleshoot.eptoeputils.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)

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
- 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

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.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.l3ext** 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

CHAPTER 7

Troubleshooting Steps for Endpoint Connectivity Problems

This chapter lists the steps for troubleshooting endpoint connectivity issues using the Cisco APIC tools, contains procedures for inspecting the operational status of your endpoints and tunnel interfaces, and explains how to connect an SFP module.

This chapter contains the following sections:

Troubleshooting Endpoint Connectivity

Step 1 Inspect the operational status of each endpoint.
The operational status will reveal any fault or misconfiguration of the endpoints. See

Inspecting the Endpoint Status, on page 108 .

Step 2 Inspect the status of the tunnel interface.
The operational status will reveal any fault or misconfiguration of the tunnel. See Inspecting the Tunnel Interface Status, on page 109.

Step 3 Perform a traceroute between the endpoint groups (EPGs).
A traceroute will reveal any problems with intermediate nodes, such as spine nodes, between the endpoints. See Performing a Traceroute Between Endpoints, on page 55.

Step 4 Configure an atomic counter on an endpoint.
The atomic counter will confirm whether the source endpoint is transmitting packets or the destination endpoint is receiving packets, and whether the number of packets received equals the number of packets sent. See Configuring Atomic Counters, on page 30.

Step 5 Inspect the contracts under each EPG.

Inspect the contracts under each EPG to make sure they allow the traffic that should flow between the EPGs. As a test, you can temporarily open the contracts to allow unrestricted traffic.

Step 6 Configure a SPAN policy to forward source packets to a monitoring node.
A packet analyzer on the monitoring node will reveal any packet issues such as an incorrect address or protocol. See Configuring a SPAN Session, on page 45.

Inspecting Endpoint and Tunnel Interface Status

This section explains how to inspect the operational status of endpoints and tunnel interfaces. Performing these procedures enables you to reveal any fault or misconfiguration of the endpoints and tunnel interfaces.

Inspecting the Endpoint Status

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 **Application Profiles**, and expand the application profile that contains the endpoint.

Step 4 Expand **Application EPGs** and click the EPG to be inspected.

Step 5 In the **Work** pane, from the list of endpoints in the **Endpoint** table, double-click the source endpoint to open the **Client End Point** dialog box.

Step 6 In the **Client End Point** dialog box, verify the endpoint properties and click the **Operational** tab.

Step 7 In the **Operational** tab, view the health, status, and fault information.
In the **Status** table, click any items with entries, such as changes, events, or faults.

Step 8 Close the **Client End Point** dialog box.

Step 9 In the **Endpoint** table, view the **Interface** entry for the endpoint and note the node and tunnel IDs.

Step 10 Repeat this procedure for the destination endpoint.

Note Occasionally, bidirectional traffic is interrupted between IP addresses in two micro-segmented EPGs deployed behind two leaf switches in the fabric. This can occur when the IP addresses are transitioning because of a configuration change from micro-segment EPG to base EPG or vice versa, on two different leaf switches at the same time, while bidirectional traffic is running. In this case, the policy tag for each remote endpoint still points to its previous EPG.

Workaround: Manually clear the remote endpoints on the switches or wait for the remote endpoint to age out. To clear the endpoints, log on to the CLI on each switch and enter the **clear system internal epm endpoint {ip | ipv6}** *ip-address* command. Then the endpoints will be relearned with the correct policy tag.

Inspecting the Tunnel Interface Status

This procedure shows how to inspect the operational status of the tunnel interface.

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

Step 2 In the submenu bar, click **Inventory**.

Step 3 In the **Navigation** pane, expand the pod and expand the node ID of the source endpoint interface.

Step 4 Under the node, expand **Interfaces**, expand **Tunnel Interfaces**, and click the tunnel ID of the source endpoint interface.

Step 5 In the **Work** pane, verify the tunnel interface properties and click the **Operational** tab.

Step 6 In the **Operational** tab, view the health, status, and fault information.
In the **Status** table, click any items with entries, such as changes, events, or faults.

Step 7 Repeat this procedure for the destination endpoint interface.

Connecting an SFP Module

When you connect an SFP module to a new card, you need to create a link speed policy for the module to communicate with the card. Follow these steps to create a link speed policy.

Step 1 Create an interface policy to specify the link speed:

Example:

```
<fabricHIfPol name=”SpeedPol” speed=”1G”/>
```

Step 2 Reference the link speed policy within an interface policy group:

Example:

```
<infraAccPortGrp name=”myGroup”>
   <infraRsHIfPol tnFabricHIfPolName=”SpeedPol”/>
</infraAccPortGrp>
```

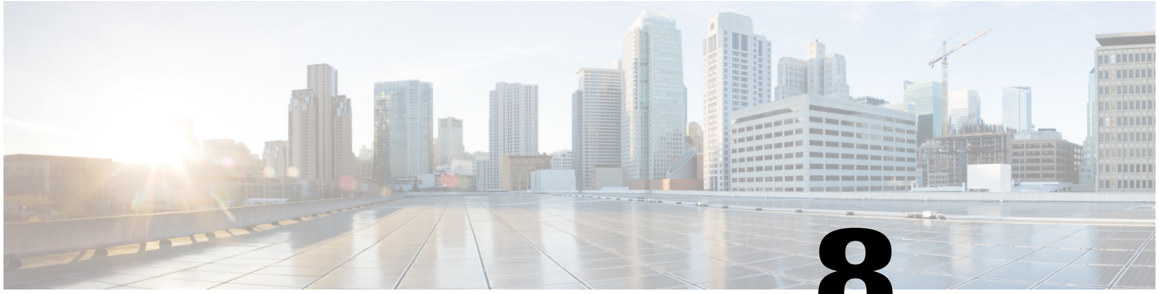

CHAPTER 8

Troubleshooting EVPN Type-2 Route Advertisement

Troubleshooting EVPN Type-2 Route Distribution to a DCIG

For optimal traffic forwarding in an EVPN topology, you can enable fabric spines to distribute host routes to a Data Center Interconnect Gateway (DCIG) using EVPN type-2 (MAC-IP) routes along with the public BD subnets in the form of BGP EVPN type-5 (IP Prefix) routes. This is enabled using the HostLeak object. If you encounter problems with route distribution, use the steps in this topic to troubleshoot.

SUMMARY STEPS

1. Verify that HostLeak object is enabled under the VRF-AF in question, by entering a command such as the following in the spine-switch CLI:
2. Verify that the config-MO has been successfully processed by BGP, by entering a command such as the following in the spine-switch CLI:
3. Verify that the public BD-subnet has been advertised to DCIG as an EVPN type-5 route:
4. Verify whether the host route advertised to the EVPN peer was an EVPN type-2 MAC-IP route:
5. Verify that the EVPN peer (a DCIG) received the correct type-2 MAC-IP route and the host route was successfully imported into the given VRF, by entering a command such as the following on the DCIG device (assuming that the DCIG is a Cisco ASR 9000 switch in the example below):

DETAILED STEPS

Step 1 Verify that HostLeak object is enabled under the VRF-AF in question, by entering a command such as the following in the spine-switch CLI:

Example:

```
spine1# ls /mit/sys/bgp/inst/dom-apple/af-ipv4-ucast/
ctrl-l2vpn-evpn  ctrl-vpnv4-ucast  hostleak  summary
```

Step 2 Verify that the config-MO has been successfully processed by BGP, by entering a command such as the following in the spine-switch CLI:

Example:

```
spine1# show bgp process vrf apple
```

Look for output similar to the following:

```
 Information for address family IPv4 Unicast in VRF apple
    Table Id                   : 0
    Table state                : UP
    Table refcount             : 3
    Peers      Active-peers    Routes     Paths      Networks   Aggregates
    0          0               0          0          0          0

    Redistribution
        None

    Wait for IGP convergence is not configured
    GOLF EVPN MAC-IP route is enabled
    EVPN network next-hop 192.41.1.1
    EVPN network route-map map_pfxleakctrl_v4
    Import route-map rtctrlmap-apple-v4
    EVPN import route-map rtctrlmap-evpn-apple-v4
```

Step 3 Verify that the public BD-subnet has been advertised to DCIG as an EVPN type-5 route:

Example:

```
spine1# show bgp l2vpn evpn 10.6.0.0 vrf overlay-1
Route Distinguisher: 192.41.1.5:4123    (L3VNI 2097154)
BGP routing table entry for [5]:[0]:[0]:[16]:[10.6.0.0]:[0.0.0.0]/224, version 2088
Paths: (1 available, best #1)
Flags: (0x000002 00000000) on xmit-list, is not in rib/evpn
Multipath: eBGP iBGP

  Advertised path-id 1
  Path type: local 0x4000008c 0x0 ref 1, path is valid, is best path
  AS-Path: NONE, path locally originated
    192.41.1.1 (metric 0) from 0.0.0.0 (192.41.1.5)
      Origin IGP, MED not set, localpref 100, weight 32768
      Received label 2097154
      Community: 1234:444
      Extcommunity:
          RT:1234:5101
          4BYTEAS-GENERIC:T:1234:444

  Path-id 1 advertised to peers:
    50.41.50.1
```

In the **Path type** entry, **ref 1** indicates that one route was sent.

Step 4 Verify whether the host route advertised to the EVPN peer was an EVPN type-2 MAC-IP route:

Example:

```
spine1# show bgp l2vpn evpn 10.6.41.1 vrf overlay-1
Route Distinguisher: 10.10.41.2:100    (L2VNI 100)
BGP routing table entry for [2]:[0]:[2097154]:[48]:[0200.0000.0002]:[32]:[10.6.41
.1]/272, version 1146
```

```
Shared RD: 192.41.1.5:4123    (L3VNI 2097154)
Paths: (1 available, best #1)
Flags: (0x00010a 00000000) on xmit-list, is not in rib/evpn
Multipath: eBGP iBGP

  Advertised path-id 1
  Path type: local 0x4000008c 0x0 ref 0, path is valid, is best path
  AS-Path: NONE, path locally originated
  EVPN network: [5]:[0]:[0]:[16]:[10.6.0.0]:[0.0.0.0] (VRF apple)
    10.10.41.2 (metric 0) from 0.0.0.0 (192.41.1.5)
      Origin IGP, MED not set, localpref 100, weight 32768
      Received label 2097154 2097154
      Extcommunity:
          RT:1234:16777216

 Path-id 1 advertised to peers:
    50.41.50.1
```

The **Shared RD** line indicates the RD/VNI shared by the EVPN type-2 route and the BD subnet.

The **EVPN Network** line shows the EVPN type-5 route of the BD-Subnet.

The **Path-id advertised to peers** indicates the path advertised to EVPN peers.

Step 5 Verify that the EVPN peer (a DCIG) received the correct type-2 MAC-IP route and the host route was successfully imported into the given VRF, by entering a command such as the following on the DCIG device (assuming that the DCIG is a Cisco ASR 9000 switch in the example below):

Example:

```
RP/0/RSP0/CPU0:asr9k#show bgp vrf apple-2887482362-8-1 10.6.41.1
Tue Sep  6 23:38:50.034 UTC
BGP routing table entry for 10.6.41.1/32, Route Distinguisher: 44.55.66.77:51
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker               2088        2088
Last Modified: Feb 21 08:30:36.850 for 28w2d
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  Local
    192.41.1.1 (metric 42) from 10.10.41.1 (192.41.1.5)
      Received Label 2097154
      Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate, imported
      Received Path ID 0, Local Path ID 1, version 2088
      Community: 1234:444
      Extended community: 0x0204:1234:444 Encapsulation Type:8 Router
MAC:0200.c029.0101 RT:1234:5101
      RIB RNH: table_id 0xe0000190, Encap 8, VNI 2097154, MAC Address: 0200.c029.0101,
IP Address: 192.41.1.1, IP table_id 0x00000000
      Source AFI: L2VPN EVPN, Source VRF: default,
Source Route Distinguisher: 192.41.1.5:4123
```

In this output, the received RD, next hop, and attributes are the same for the type-2 route and the BD subnet.

CHAPTER 9

Performing a Rebuild of the Fabric

This chapter explains how to rebuild your fabric.

Rebuilding the Fabric

Caution This procedure is extremely disruptive. It eliminates the existing fabric and recreates a new one.

This procedure allows you to rebuild (reinitialize) your fabric, which you may need to do for any of the following reasons:

- To change the TEP IPs
- To change the Infra VLAN
- To change the fabric name
- To perform TAC troubleshooting tasks

Deleting the APICs erases the configuration on them and brings them up in the startup script. Performing this on the APICs can be done in any order, but ensure that you perform the procedure on all of them (every leaf and spine in the fabric).

Before You Begin

Ensure that the following is in place:

- Regularly scheduled backups of the configuration
- Console access to the leaves and spines
- A configured and reachable CIMC, which is necessary for KVM console access

- No Java issues

Step 1 If you would like to retain your current configuration, you can perform a configuration export using the following procedure: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_Using_Import_Export_to_Recover_Config_States.html

Step 2 Erase the configuration on the APICs by connecting to the KVM console and entering the following commands:

a) **>acidiag touch clean**

b) **>acidiag touch setup**

c) **>acidiag reboot**

Ensure that each node boots up in fabric discovery mode and is not part of the previously configured fabric.

Note The **acidiag touch** command alone is not useful for this procedure, because it does not bring the APIC up in the startup script.

Caution It is extremely important that you ensure that all previous fabric configurations have been removed. If any previous fabric configuration exists on even a single node, the fabric cannot be rebuilt.

Step 3 When all previous configurations have been removed, run the startup script for all APICs. At this point, you can change any of the above values, TEP, TEP Vlan, and/or Fabric Name. Ensure that these are consistent across all APICs. For more information, refer to: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/getting-started/b_APIC_Getting_Started_Guide/b_APIC_Getting_Started_Guide_chapter_01.html#concept_F46E2193E3134CD090B65B16038D11A9.

Step 4 Log in to apic1 and perform a configuration import using the following procedure: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_Using_Import_Export_to_Recover_Config_States.html.

Step 5 Wait for a few minutes as the fabric now uses the previous fabric registration policies to rebuild the fabric over the nodes. (Depending on the size of the fabric, this step may take awhile.)

CHAPTER 10

Verifying IP-Based EPG Configurations

There are two types of endpoint groups (EPGs) that you can create: application EPGs and IP-based EPGs. IP-based EPGs differ from regular application EPGs in that they are microsegment EPGs. This chapter explains how to verify that your IP-based EPG configurations are properly classified as IP-based using the GUI or using switch commands.

This chapter contains the following sections:

Verifying IP-Based EPG Configurations Using the GUI

This procedure explains how to verify that you have correctly configured an IP-based EPG using the GUI and Visore tool.

Step 1 Verify that the IP-based EPG you created is listed under the **uSeg EPGs** folder in the GUI (shown in the following screen capture).
Note that there is one IP-based EPG listed under uSeg EPGs named "IP" that was created using the REST API.

Step 2 Verify that the information is correct in the EPG - IP properties screen (right side window pane) for each EPG IP (IP-based EPG), shown as follows:

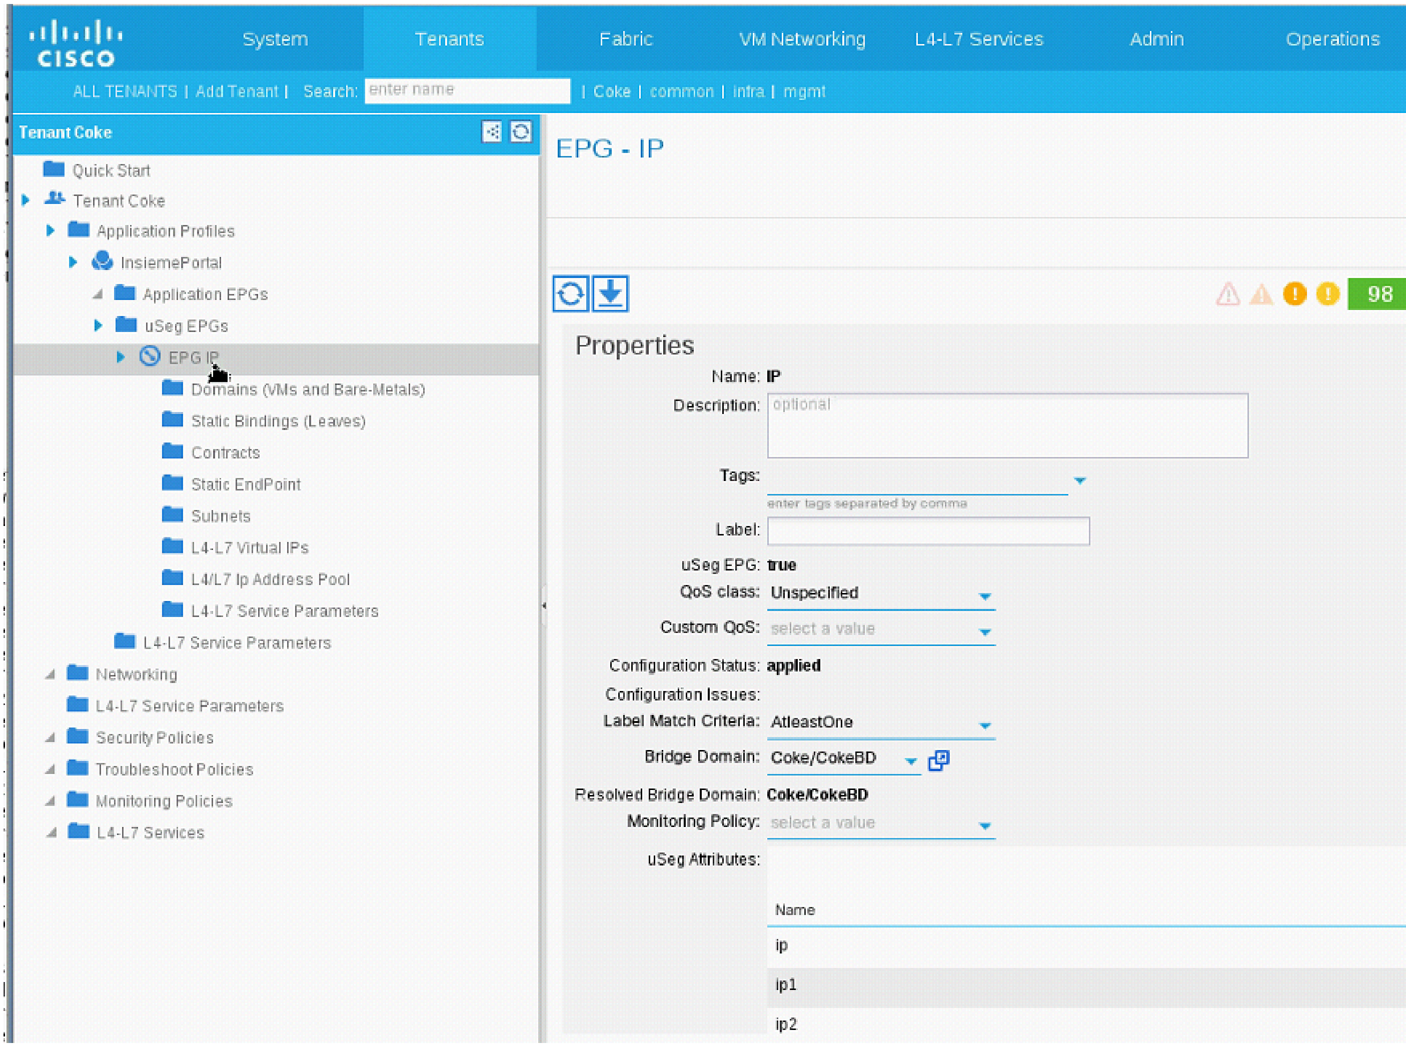

Properties

Name: IP

Description: optional

Tags: enter tags separated by comma

Label:

uSeg EPG: true

QoS class: Unspecified

Custom QoS: select a value

Configuration Status: applied

Configuration Issues:

Label Match Criteria: AtleastOne

Bridge Domain: Coke/CokeBD

Resolved Bridge Domain: Coke/CokeBD

Monitoring Policy: select a value

uSeg Attributes:

Name
ip
ip1
ip2

Note the list of IP-based EPGs and IP addresses that are shown at the bottom of the screen.

Step 3 From your web browser, enter the APIC IP address followed by "/visore.html" (shown as follows). Visore is a tool that allows you to view all the objects in the system, such as EPGs. You can use Visore to verify that your IP-based EPGs have been properly configured. For more information about Visore, see the *Application Policy Infrastructure Controller Visore Tool Introduction* document.

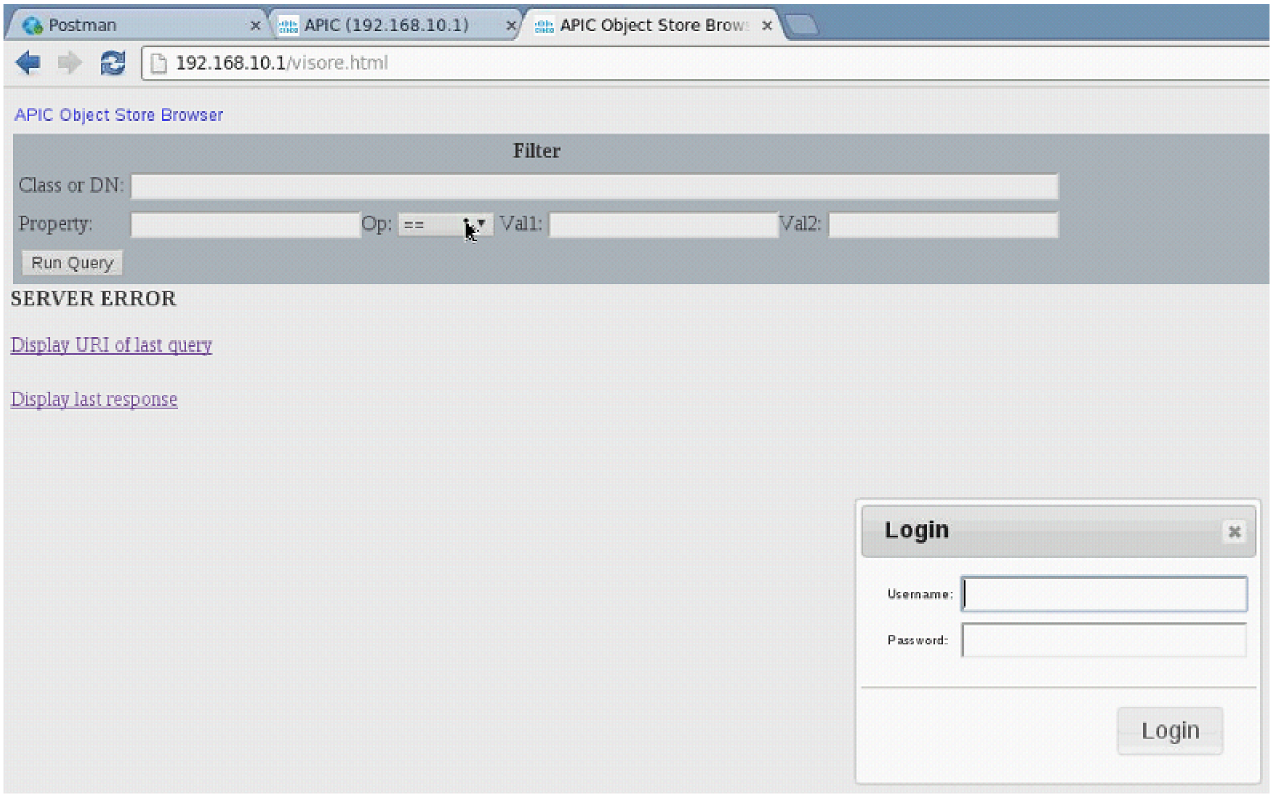

Step 4 Enter your username and password then click **Login** to log into Visore.

Step 5 Run a query for the IP-based EPGs that you verified in the GUI by entering the name of the class in the field next to **Class or DN** (for example, "fvAEPg"), shown as follows:

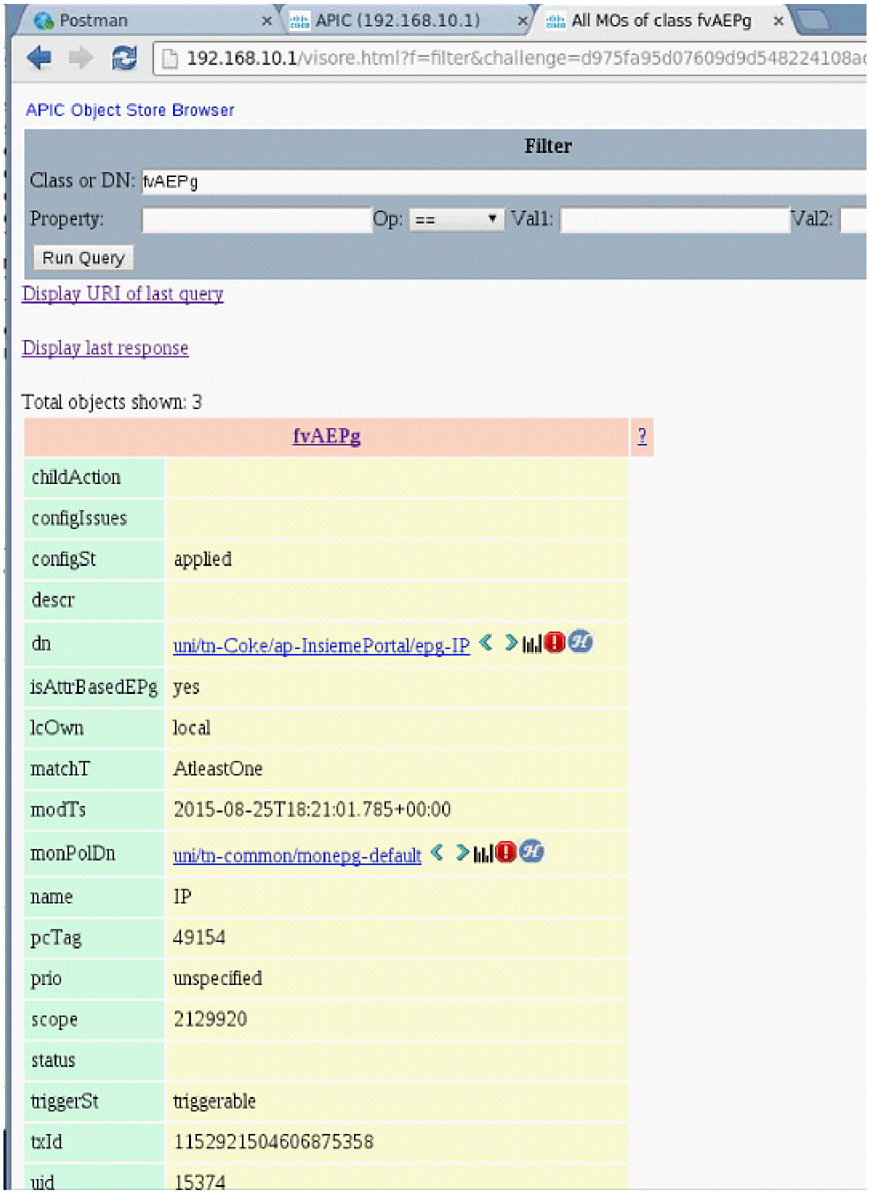
Postman
APIC (192.168.10.1)
All MOs of class fvAEPg
192.168.10.1/visore.html?f=filter&challenge=d975fa95d07609d9d548224108a
APIC Object Store Browser
Filter
Run Query
Display URI of last query
Display last response

Class or DN: fvAEPg

Property: Op: == Val1: Val2:

Total objects shown: 3

fvAEPg		?
childAction		
configIssues		
configSt	applied	
descr		
dn	uni/tn-Coke/ap-InsiemePortal/epg-IP	
isAttrBasedEPg	yes	
lcOwn	local	
matchT	AtleastOne	
modTs	2015-08-25T18:21:01.785+00:00	
monPolDn	uni/tn-common/monepg-default	
name	IP	
pcTag	49154	
prio	unspecified	
scope	2129920	
status		
triggerSt	triggerable	
txId	1152921504606875358	
uid	15374	

Note This is a view from the APIC point of view. You can see that the "Total objects shown" above is "3", meaning there are three EPGs that were downloaded to the switch. You can see that the IP-based EPG that was previously listed in the GUI as "IP" is now shown next to "dn". Also note that "yes" is displayed next to "isAttrBasedEPg", which means that this has been properly configured as an IP-based EPG. You can verify all the objects have been configured successfully using Visore, including both application EPGs and IP-based EPGs.

Step 6 This is a view from the switch point of view. On the switch, you can run a query for the fvEpP class to see the EPGs and check for the "crtrnEnabled" attribute. It will be set to "yes" for IP-based EPGs.

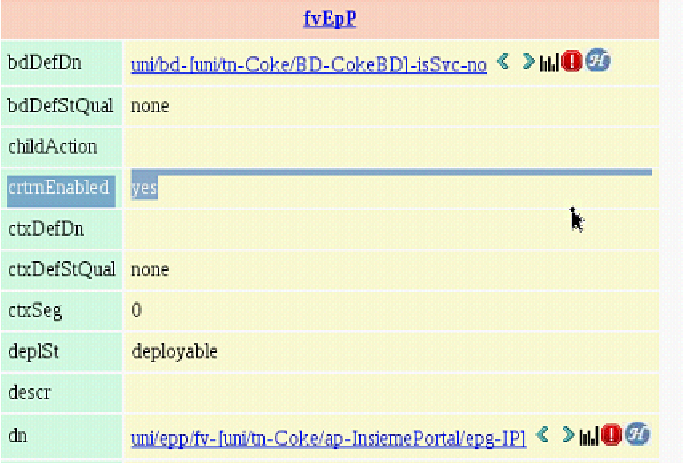

fvEpP	
bdDefDn	uni/bd-[uni/tn-Coke/BD-CokeBD]-isSvc-no
bdDefStQual	none
childAction	
crtrnEnabled	yes
ctxDefDn	
ctxDefStQual	none
ctxSeg	0
deplSt	deployable
descr	
dn	uni/epp/fv-[uni/tn-Coke/ap-InsiemePortal/epg-IP]

Step 7 Verify that under this EPG, the children of the EPG are shown with IP addresses to ensure a proper configuration. For each IP address configured, there is one object (named "l3IpCktEp") that the switch uses to classify the traffic, shown as follows:

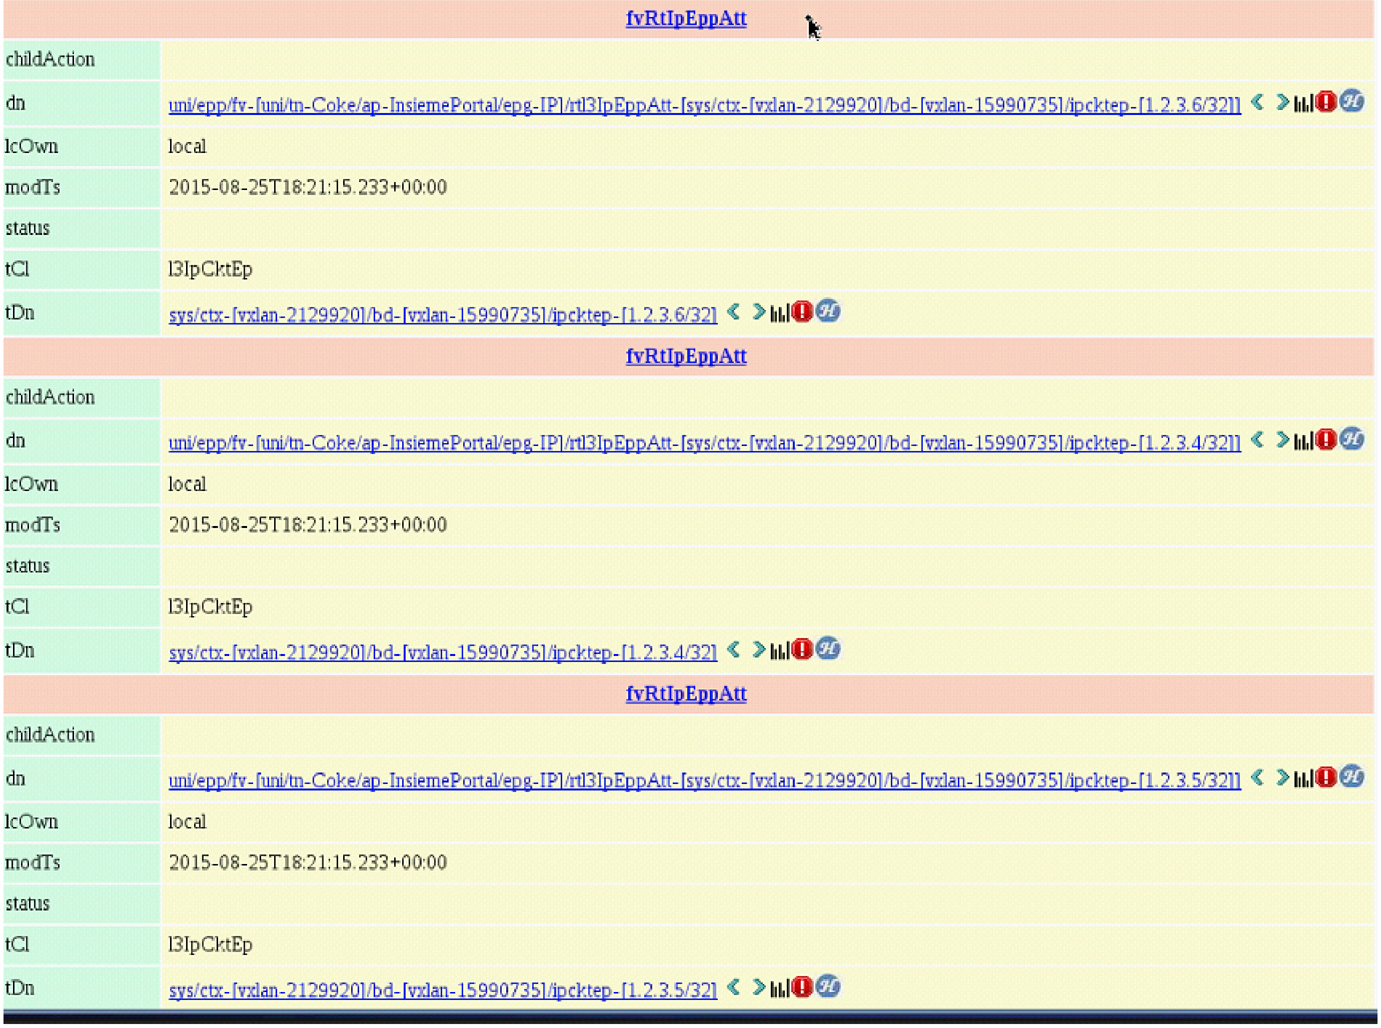

fvRtIpEppAtt	
childAction	
dn	uni/epp/fv-[uni/tn-Coke/ap-InsiemePortal/epg-IP]/rtl3IpEppAtt-[sys/ctx-[vxlan-2129920]/bd-[vxlan-15990735]/ipcktep-[1.2.3.6/32]]
lcOwn	local
modTs	2015-08-25T18:21:15.233+00:00
status	
tCl	l3IpCktEp
tDn	sys/ctx-[vxlan-2129920]/bd-[vxlan-15990735]/ipcktep-[1.2.3.6/32]

fvRtIpEppAtt	
childAction	
dn	uni/epp/fv-[uni/tn-Coke/ap-InsiemePortal/epg-IP]/rtl3IpEppAtt-[sys/ctx-[vxlan-2129920]/bd-[vxlan-15990735]/ipcktep-[1.2.3.4/32]]
lcOwn	local
modTs	2015-08-25T18:21:15.233+00:00
status	
tCl	l3IpCktEp
tDn	sys/ctx-[vxlan-2129920]/bd-[vxlan-15990735]/ipcktep-[1.2.3.4/32]

fvRtIpEppAtt	
childAction	
dn	uni/epp/fv-[uni/tn-Coke/ap-InsiemePortal/epg-IP]/rtl3IpEppAtt-[sys/ctx-[vxlan-2129920]/bd-[vxlan-15990735]/ipcktep-[1.2.3.5/32]]
lcOwn	local
modTs	2015-08-25T18:21:15.233+00:00
status	
tCl	l3IpCktEp
tDn	sys/ctx-[vxlan-2129920]/bd-[vxlan-15990735]/ipcktep-[1.2.3.5/32]

Once the configuration is there, when the packets arrive, the switch uses these objects to classify them.

Step 8 Verify that the pcTags for all the endpoints and IP addresses that you configured match. Every EPG has a pcTag. All the endpoints that match with the IP addresses you configured are classified into this pcTag. Every endpoint has an IP address that you can run a class query on. When you are troubleshooting, you want to verify whether these endpoints (servers) are properly getting classified into this IP-based EPG or not. (The pcTags should match for the IP-based EPG.)

Verifying IP-EPG Configurations Using Switch Commands

This procedure explains how to use switch commands to verify you IP-EPG ("IpCkt") configurations.

Step 1 Log in to the leaf.

Step 2 Navigate to the /mit/sys directory.

Step 3 In the /mit/sys directory, find ctx (vrf context directory)

Step 4 In the VRF cts directory, go to the specific BD directory where the IpCkt is configured.
You should see the IpCkt.

Note "IpCkt" and "IP-EPG" are used interchangeably in this document.

Step 5 Navigate to the directory and the "cat summary" gives you the information regarding IpCkt.

Step 6 Ensure that the summary's "operSt' does not say "unsupported".

Step 7 Find out the VLAN ID that corresponds to the BD where the IpCkt is configured.

Note The VLAN ID can be found through any of the **show vlan internal bd-info** commands or through the **show system internal epm vlan all** command.

Step 8 Once you find the VLAN ID of the BD, issue **show system internal epm <vlan-id> detail**.
Here you should be able to see all the configured IpCkts with a specific sclass. (It should match that of what you see in the /mit/sys directory.)

Step 9 Repeat the steps for vsh_lc that you followed for vsh.

Step 10 Send the traffic with an IP matching the IpCtk in the BD, and through **show system internal epm endp ip <a.b.c.d>**, you can verify that the learned IP has the IP-flags for "sclass" and a specific sclass value.

Step 11 Repeat the steps for vsh_lc that you followed for vsh.

List of the Switch Troubleshooting Commands Used in this Procedure:

```
Cd /mits/sys/ctx-vxlan…/bd-vxlan…
     - cat summary
Vsh -c “show system internal epm vlan all” or
Vsh -c “show vlan internal bd-info”
Vsh -c “show system internal epm vlan <vlan-id> detail”
Vsh -c “show system internal epm endp ip <a.b.c.d>"
Vsh_lc -c “show system internal epm vlan all” or
Vsh_lc -c “show vlan internal bd-info”
Vsh_lc -c “show system internal epm vlan <vlan-id> detail”
vsh_lc -c “show system internal epm endp ip <a.b.c.d>”
vsh_lc -c “show system internal epm epg”
```

CHAPTER 11

Recovering a Disconnected Leaf

If all fabric interfaces on a leaf are disabled (interfaces connecting a leaf to the spine) due to a configuration pushed to the leaf, connectivity to the leaf is lost forever and the leaf becomes inactive in the fabric. Trying to push a configuration to the leaf does not work because connectivity has been lost. This chapter describes how to recover a disconnected leaf.

Recovering a Disconnected Leaf Using the REST API

To recover a disconnected leaf, at least one of the fabric interfaces must be enabled using the following process. The remaining interfaces can be enabled using the GUI, REST API, or CLI.

To enable the first interface, post a policy using the REST API to delete the policy posted and bring the fabric ports Out-of-Service. You can post a policy to the leaf to bring the port that is Out-of-Service to In-Service as follows:

Note In the following examples, the assumption is that 1/49 is one of the leaf ports connecting to the spine.

Step 1 Clear the blacklist policy from the APIC (using the REST API).

Example:

```
$APIC_Address/api/policymgr/mo/.xml
<polUni>
    <fabricInst>
        <fabricOOServicePol>
                       <fabricRsOosPath tDn="topology/pod-1/paths-$LEAF_Id/pathep-[eth1/49]"
lc="blacklist" status ="deleted" />
        </fabricOOServicePol>
    </fabricInst>
</polUni>
```

Step 2 Post a local task to the node itself to bring up the interfaces you want using **l1EthIfSetInServiceLTask**.

Example:

```
$LEAF_Address/api/node/mo/topology/pod-1/node-$LEAF_Id/sys/action.xml
<actionLSubj oDn="sys/phys-[eth1/49]">
<l1EthIfSetInServiceLTask adminSt='start'/>
</actionLSubj>
```

CHAPTER 12

Determining Why a PIM Interface Was Not Created

A PIM interface (pim:if) is created for L3Out interfaces (note that L3Out SVI interfaces are not supported), multicast tunnel interfaces (per VRF), SVI interfaces corresponding to PIM-enabled pervasive BDs, and loopback interfaces on border leafs (each per VRF).

This chapter contains troubleshooting information for situations where the pim:if is not being created. For more information on PIM, see the *Cisco ACI and Layer 3 Multicast with Cisco ACI* and the *Cisco Application Centric Infrastructure Fundamentals* guides.

This chapter contains the following sections:

A PIM Interface Was Not Created For an L3Out Interface

If a PIM interface (pim:If) is not being created for an L3Out interface, confirm the following:

1. PIM is enabled on the L3Out. If PIM is disabled, enable it.
2. If PIM is enabled on the container L3Out, confirm that a multicast l3ext:InstP has been created with "__int_" as a prefixed name. This multicast l3ext:InstP is used to deploy L3Out PIM policies to the switches. There should be one multicast l3ext:InstP per L3Out.

Note

- If a multicast l3ext:InstP exists on the IFC, we can check whether a corresponding fv:RtdEpP is created and deployed on each switch where there is an interface in that L3Out.
- We do not support an L3Out SVI interface for PIM.

A PIM Interface Was Not Created For a Multicast Tunnel Interface

If a PIM interface (pim:if) is not created for a multicast tunnel interface (tunnel:If), confirm the following:

1. The corresponding tunnel:If has been created.

Note The tunnel:If should have type “underlay-mcast."

2. Each mcast-enabled VRF has created an mcast tunnel.
3. The destination IP field of the tunnel:If is populated with a valid GIPO address.
4. If the tunnel:If is not populated with a valid GIPO address, check the pim:CtxP on the IFC and the pim:CtxDef on the switches to make sure GIPO is allocated correctly.
5. The source IP of the tunnel:If has the loopback address of an L3Out for BL and “127.0.0.100” for NBL.

A PIM Interface Was Not Created For a Multicast-Enabled Bridge Domain

If a PIM interface (pim:if) is not created for a multicast-enabled bridge domain (BD), confirm the following:

1. The corresponding BD or corresponding Ctx has PIM enabled.
2. The corresponding BD is pervasive.
3. The pervasive BD-based pim:If takes default parameters.

Note For interaction with igmp snooping, when PIM is enabled on a pervasive BD, the routing bit should be automatically enabled for the corresponding igmpsnoop:If.

CHAPTER 13

Confirming the Port Security Installation

This chapter explains how to confirm the port security installation in the APIC and leaf switch using Visore and how to confirm port security has been programmed in the hardware using the Cisco NX-OS-style CLI. For information about configuring port security, see the *Cisco Port Security* document.

This chapter contains the following sections:

Confirming Your Port Security Installation Using Visore

Step 1 On the Cisco APIC, run a query for the l2PortSecurityPol class in Visore to verify the port security policy installation.

Step 2 On the leaf switch, run a query for l2PortSecurityPolDef in Visore to confirm that the concrete object exists on the interface.
If you have confirmed that port security is installed on the Cisco APIC and leaf switch, use the Cisco NX-OS CLI to confirm that port security has been programmed in the hardware.

Confirming Your Hardware Port Security Installation Using the Cisco NX-OS CLI

Step 1 View the port security status on the switch interface as follows:

Example:

```
switch# show system internal epm interface ethernet 1/35 det
name : Ethernet1/35 ::: if index : 0x1a022000 ::: state : UP
vPC : No ::: EPT : 0x0
```

```
MAC Limit : 8 ::: Learn Disable : No ::: PortSecurity Action : Protect
VLANs : 4-23
Endpoint count : 5
Active Endpoint count : 5
switch# show system internal epm interface port-channel 1 det

name : port-channel1 ::: if index : 0x16000000 ::: state : UP
vPC : No ::: EPT : 0x0
MAC Limit : 6 ::: Learn Disable : No ::: PortSecurity Action : Protect
VLANs :
Endpoint count : 0
Active Endpoint count : 0
Number of member ports : 1
Interface : Ethernet1/34    /0x1a021000
::::
```

Step 2 View the port security status on the module interface as follows:

Example:

```
module-1# show system internal epmc interface ethernet 1/35 det
if index : 0x1a022000 ::: name : Ethernet1/35 ::: tun_ip = 0.0.0.0
MAC limit : 8 ::: is_learn_disable : No ::: MAC limit action: Protect
pc if index : 0 ::: name :
is_vpc_fc FALSE  ::: num_mem_ports : 0
 interface state : up
Endpoint count : 5
EPT : 0

module-1# show system internal epmc interface port-channel 1 det
if index : 0x16000000 ::: name : port-channel1 ::: tun_ip = 0.0.0.0
MAC limit :  6 ::: is_learn_disable : No ::: MAC limit action: Protect
pc if index : 0 ::: name :
is_vpc_fc FALSE  ::: num_mem_ports : 1
 interface state : up
Endpoint count : 0
EPT : 0
::::
```

Step 3 View the port security status on the leaf switch as follows:

Example:

```
swtb15-leaf2# show system internal epm interface ethernet 1/35 det

name : Ethernet1/35 ::: if index : 0x1a022000 ::: state : UP
vPC : No ::: EPT : 0x0
MAC Limit : 5 ::: Learn Disable : Yes ::: PortSecurity Action : Protect
VLANs : 4-23
Endpoint count : 5
Active Endpoint count : 5
::::
```

Step 4 Confirm the MAC limit on the module interface as follows:

Example:

```
module-1# show system internal eltmc info interface port-channel1 | grep mac_limit
   mac_limit_reached:              0   :::       mac_limit:              8
port_sec_feature_set:              1   ::: mac_limit_action:              1
```

Example:

```
module-1# show system internal eltmc info interface ethernet 1/35 | grep mac_limit
   mac_limit_reached:              0   :::       mac_limit:              8
port_sec_feature_set:              1   ::: mac_limit_action:              1
```

Step 5 View the port security status in the module and confirm the MAC limit as follows:

Example:

```
module-1#  show system internal epmc interface ethernet 1/35 det
if index : 0x1a022000 ::: name : Ethernet1/35 ::: tun_ip = 0.0.0.0
MAC limit : 5 ::: is_learn_disable : Yes ::: MAC limit action: Protect
pc if index : 0 ::: name :
is_vpc_fc FALSE  ::: num_mem_ports : 0
 interface state : up
Endpoint count : 5
EPT : 0
::::
```

Example:

```
module-1# show system internal eltmc info interface ethernet 1/35 | grep mac_limit
   mac_limit_reached:              1   :::       mac_limit:              5
port_sec_feature_set:              1   ::: mac_limit_action:              1
module-1# exit
```

CHAPTER 14

Troubleshooting QoS Policies

This section provides solutions for troubleshooting QoS policies.

Troubleshooting Cisco APIC QoS Policies

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

Problem	Solution
Unable to update a configured QoS policy.	**1** Invoke the following API to ensure that `qospDscpRule` is present on the leaf. `GET https://192.0.20.123/api/node/class/qospDscpRule.xml` **2** Ensure that the QoS rules are accurately configured and associated to the EPG ID to which the policy is attached. Use the following commands to verify the configuration. `leaf1#` **show vlan** `leaf1#` **show system internal aclqos qos policy detail** `apic1#` **show running-config tenant** *tenant-name* **policy-map type qos** *custom-qos-policy-name* `apic1#` **show running-config tenant** *tenant-name* **application** *application-name* **epg** *epg-name*

CHAPTER 15

Determining the Supported SSL Ciphers

This chapter explains how to determine which SSL ciphers are supported.

About SSL Ciphers

The Cisco Application Centric Infrastructure (ACI) Representational State Transfer (REST) Application Programming Interface (API) has gone through an evolution from the day the solution debuted to recent versions where the HTTPS/SSL/TLS support has gotten increasingly more stringent. This document is intended to cover the evolution of HTTPS, SSL, and TLS support on the Cisco ACI REST API and provide customers with a guide of what is required for a client to utilize the REST API securely.

HTTPS is a protocol that utilizes either Secure Socket Layers (SSL) or Transport Layer Security (TLS) to form a secure connection for a HTTP session. SSL or TLS is used to encrypt the traffic between a client and a HTTP server. In addition, servers that support HTTPS have a certificate that can usually be used by the client to verify the server's authenticity. This is the opposite of the client authenticating with the server. In this case, the server is saying, "I am server_xyz and here is the certificate that proves it." The client can then utilize that certificate to verify the server is "server_xyz."

There are other important aspects to SSL/TLS that involve the supported encryption ciphers available in each protocol as well as the inherent security of the SSL or TLS protocols. SSL has gone through three iterations - SSLv1, SSLv2 and SSLv3 - all of which are now considered insecure. TLS has gone through three iterations - TLSv1, TLSv1.1 and TLSv1.2 - of which only TLSv1.1 and TLSv1.2 are considered "secure." Ideally, a client should utilize the highest available TLS version it can and the server should support only TLSv1.1 and TLSv1.2. However, most servers must keep TLSv1 for outdated clients.

Almost all modern browsers support both TLSv1.1 and TLSv1.2. However, a client that utilizes HTTPS may not be a browser. The client may be a java application or a python script that communicates with a web server and must negotiate HTTPS/TLS. In this type of a situation, the questions of what is supported and where becomes much more important.

Determining the Supported SSL Ciphers Using the CLI

Before You Begin

This section describes how to use the CLI to determine which SSL ciphers are supported.

Step 1 Get the supported ciphers in your openssl environment, shown as follows:

Example:

```
openssl ciphers 'ALL:eNULL'
```

Step 2 Separate the ciphers using sed or some other tool, shown as follows:

Example:

```
openssl ciphers 'ALL:eNULL' |  sed -e 's/:/\n/g'
```

Step 3 Loop over the ciphers and poll the APIC to see which ones are supported, shown as follows:

Example:

```
openssl s_client -cipher ?<some cipher to test>? -connect <apic ipaddress>:<ssl port, usually 443>
```

See the following example cipher:

Example:

```
openssl s_client -cipher ?ECDHE-ECDSA-AES128-GCM-SHA256? -connect 10.1.1.14:443
```

Note If the response contains `CONNECTED`, then the cipher is supported.

CHAPTER 16

Removing Unwanted _ui_ Objects

Caution

Changes made through the APIC Basic GUI can be seen, but cannot be modified in the Advanced GUI, and changes made in the Advanced GUI cannot be rendered in the Basic GUI. The Basic GUI is kept synchronized with the NX-OS style CLI, so that if you make a change from the NX-OS style CLI, these changes are rendered in the Basic GUI, and changes made in the Basic GUI are rendered in the NX-OS style CLI, but the same synchronization does not occur between the Advanced GUI and the NX-OS style CLI.

Cisco recommends that you do not mix configuration modes (Advanced or Basic). When you make a configuration in either mode and change the configuration using the other mode, unintended changes can occur. For example, if you apply an interface policy to two ports using Advanced mode and then change the settings of one port using Basic mode, your changes might be applied to both ports.

If you make changes with the Basic GUI or the NX-OS CLI before using the Advanced GUI, this may also inadvertantly cause objects to be created (with names prepended with `_ui_`) which cannot be changed or deleted in the Advanced GUI.

If you make changes with the Basic GUI or the NX-OS CLI before using the Advanced GUI, this may inadvertently cause objects to be created (with names prepended with `_ui_`) which cannot be changed or deleted in the Advanced GUI.

For the steps to remove such objects, see Removing Unwanted _ui_ Objects Using the REST API, on page 137.

Removing Unwanted _ui_ Objects Using the REST API

If you make changes with the Basic GUI or the NX-OS CLI before using the Advanced GUI, and objects appear in the Advanced GUI (with names prepended with `_ui_`), these objects can be removed by performing a REST API request to the API, containing the following:

- The Class name, for example **infraAccPortGrp**
- The Dn attribute, for example **dn="uni/infra/funcprof/accportgrp-__ui_l101_eth1--31"**
- The Status attribute set to **status="deleted"**

Perform the POST to the API with the following steps:

Step 1 Log on to a user account with write access to the object to be removed.

Step 2 Send a POST to the API such as the following example:

```
POST https://192.168.20.123/api/mo/uni.xml
Payload:<infraAccPortGrp dn="uni/infra/funcprof/accportgrp-__ui_l101_eth1--31" status="deleted"/>
```

APPENDIX A

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: • 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. **Note** Peer chassis IDs can be incorrect for APICs not currently in the cluster.
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.

Option	Function
rvread *service*	Displays the data layer state for a service on all shards across all replicas. **Note** For an example, see Examples, on page 144
rvread *service shard*	Displays the data layer state for a service on a specific shard across all replicas. **Note** For an example, see Examples, on page 144
rvread *service shard replica*	Displays the data layer state for a service on a specific shard and replica. **Note** For an example, see Examples, on page 144
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.

Table 2: Service IDs

Service	ID
cliD	1
controller	2
eventmgr	3
extXMLApi	4
policyelem	5
policymgr	6
reader	7
ae	8
topomgr	9
observer	10
dbgr	11

Service	ID
observerelem	12
dbgrelem	13
vmmmgr	14
nxosmock	15
bootmgr	16
appliancedirector	17
adrelay	18
ospaagent	19
vleafelem	20
dhcpd	21
scripthandler	22
idmgr	23
ospaelem	24
osh	25
opflexagent	26
opflexelem	27
confelem	28
vtap	29

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.
imageurl	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. • 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.

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:

```
admin@apic1:~> acidiag version
1.0.0.414

admin@apic1:~> acidiag verifyapic
openssl_check: certificate details
subject= CN=Insieme,O=Insieme Networks,L=SanJose,ST=CA,C=US
issuer= O=Default Company Ltd,L=Default City,C=XX
notBefore=Jul 19 20:40:32 2013 GMT
notAfter=Jul 19 20:40:32 2014 GMT
openssl_check: passed
ssh_check: passed
all_checks: passed

admin@apic1:~> acidiag avread
Local appliance ID=1 ADDRESS=10.0.0.1 TEP ADDRESS=10.0.0.0/16
CHASSIS_ID=10220833-ea00-3bb3-93b2-ef1e7e645889
Cluster of 3 lm(t):1(2014-07-12T19:54:04.877+00:00) appliances
  (out of targeted 3 lm(t):3(2014-07-12T19:55:03.442+00:00))
   with FABRIC_DOMAIN name=mininet set to version=1.0(0.414)
lm(t):3(2014-07-12T19:55:13.564+00:00)
    appliance id=1 last mutated at 2014-07-12T19:46:06.831+00:00 address=10.0.0.1 tep
address=10.0.0.0/16
       oob address=192.168.10.1/24 version=1.0(0.414) lm(t):1(2014-07-12T19:54:05.146+00:00)

       chassisId=10220833-ea00-3bb3-93b2-ef1e7e645889 lm(t):1(2014-07-12T19:54:05.146+00:00)

        commissioned=1 registered=1 active=yes(zeroTime)
        health=(applnc:255 lm(t):1(2014-07-12T20:01:22.934+00:00) svc's)
    appliance id=2 last mutated at 2014-07-12T19:51:10.649+00:00 address=10.0.0.2 tep
address=10.0.0.0/16
```

```
       oob address=192.168.10.2/24 version=1.0(0.414) lm(t):2(2014-07-12T19:54:05.064+00:00)

       chassisId=5d74122c-2ab9-3ccb-b06d-f620d5e20ccd lm(t):2(2014-07-12T19:54:05.064+00:00)

        commissioned=1 registered=1 active=yes(2014-07-12T19:51:10.651+00:00)
        health=(applnc:255 lm(t):2(2014-07-12T20:01:22.442+00:00) svc's)
    appliance id=3 last mutated at 2014-07-12T19:54:05.028+00:00 address=10.0.0.3 tep
address=10.0.0.0/16
       oob address=192.168.10.3/24 version=1.0(0.414) lm(t):3(2014-07-12T19:54:05.361+00:00)

       chassisId=71355d49-6fe7-3a78-a361-72d6c1e3360c lm(t):3(2014-07-12T19:54:05.361+00:00)

        commissioned=1 registered=1 active=yes(2014-07-12T19:54:05.029+00:00)
        health=(applnc:255 lm(t):3(2014-07-12T20:01:22.892+00:00) svc's)
clusterTime=<diff=0 common=2014-07-14T16:52:20.343+00:00 local=2014-07-14T16:52:20.343+00:00
 pF=<displForm=0
    offsSt=0 offsVlu=0 lm(t):3(2014-07-12T19:55:03.750+00:00)>>
---------------------------------------------

admin@apic1:~> 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)>>

admin@apic1:~> 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)>>
```

APPENDIX B

Configuring Export Policies for Troubleshooting

Export policies enable you 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.

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 for exported files cannot be an IPv6 address.

Note Do not trigger **Tech Support** from more than five nodes simultaneously, especially if they are to be exported into the APIC or to an external server with insufficient bandwidth and compute resources.

In order to collect **Tech Support** from all 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).

Configuring a Remote Location

Configuring a Remote Location Using the GUI

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

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 a Remote Location Using the REST API

This procedure explains how to create a remote location using the REST API.

```
<fileRemotePath name="local" host="host or ip" protocol="ftp|scp|sftp" remotePath="path to
folder" userName="uname" userPasswd="pwd" />
```

Configuring a Remote Location Using the NX-OS Style CLI

In the ACI fabric, you can configure one or more remote destinations for exporting techsupport or configuration files.

SUMMARY STEPS

1. **configure**
2. **[no] remote path** *remote-path-name*
3. **user** *username*
4. **path** {**ftp** | **scp** | **sftp**} *host*[*:port*] [**remote-directory**]

DETAILED STEPS

	Command or Action	Purpose
Step 1	**configure** **Example:** `apic1# configure`	Enters global configuration mode.
Step 2	**[no] remote path** *remote-path-name* **Example:** `apic1(config)# remote path myFiles`	Enters configuration mode for a remote path.
Step 3	**user** *username* **Example:** `apic1(config-remote)# user admin5`	Sets the user name for logging in to the remote server. You are prompted for a password.
Step 4	**path** {**ftp** | **scp** | **sftp**} *host*[*:port*] [**remote-directory**] **Example:** `apic1(config-remote)# path sftp filehost.example.com:21 remote-directory /reports/apic`	Sets the path and protocol to the remote server. You are prompted for a password.

Examples

This example shows how to configure a remote path for exporting files.

```
apic1# configure
apic1(config)# remote path myFiles
apic1(config-remote)# user admin5
You must reset the password when modifying the path:
Password:
Retype password:
apic1(config-remote)# path sftp filehost.example.com:21 remote-directory /reports/apic
You must reset the password when modifying the path:
Password:
Retype password:
```

Sending an On-Demand Tech Support File

Sending an On-Demand Techsupport File Using the GUI

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 TechSupport** and choose **Create On-demand TechSupport**.

Step 5 In the **Create On-demand TechSupport** dialog box, perform the following actions:

a) In the **Name** field, enter a name for the techsupport file export policy.

b) To export the file to the controller instead of a remote destination, choose **Export to Controller**.

c) From the **Export Destination** drop-down list, choose the profile of the destination host that will receive the techsupport file.
If no profile appears for the desired destination, you can choose **Create Remote Location** to define it now.

d) From the **Data Container** drop-down list, choose **uni/fabric/tscont**.

e) If the desired source device (leaf or spine) does not appear In the **Source Nodes** table, click the + icon, choose a device, and click **Update**.

f) In the **Source Nodes** table, double-click the source name and click the blue icon to the right of the drop-down list to open the **System Information** window for the source device.
Use the tabs to examine the information of the source device.

g) In the **State** field, click the **triggered** radio button to enable sending of the file.

h) Click **Submit** to send the techsupport 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 TechSupport policy in the **Navigation** pane, then click the **OPERATIONAL** tab in the **Work** pane. The controller is displayed in the **EXPORT LOCATION** field.

i) Right-click the policy name and choose **Collect Tech Support**.

j) Choose **Yes** to begin collecting tech support information.

Sending an On-Demand TechSupport File Using the REST API

Step 1 Set the remote destination for a technical support file using the REST API, by sending a POST with XML such as the following example:

Example:

```
<fileRemotePath userName="" remotePort="22" remotePath="" protocol="sftp" name="ToSupport"
host="192.168.200.2"
dn="uni/fabric/path-ToSupport" descr="">

<fileRsARemoteHostToEpg tDn="uni/tn-mgmt/mgmtp-default/oob-default"/>
```

```
</fileRemotePath>
```

Step 2 Generate an on-demand technical support file using the REST API by sending a POST with XML such as the following:

Example:

```
<dbgexpTechSupOnD upgradeLogs="no" startTime="unspecified" name="Tech_Support_9-20-16"
exportToController="no"
endTime="unspecified" dn="uni/fabric/tsod-Tech_Support_9-20-16" descr="" compression="gzip"
category="forwarding" adminSt="untriggered">

<dbgexpRsExportDest tDn="uni/fabric/path-ToSupport"/>

<dbgexpRsTsSrc tDn="topology/pod-1/node-102/sys"/>

<dbgexpRsTsSrc tDn="topology/pod-1/node-103/sys"/>

<dbgexpRsTsSrc tDn="topology/pod-1/node-101/sys"/>

<dbgexpRsData tDn="uni/fabric/tscont"/>

</dbgexpTechSupOnD>
```

APPENDIX C

Finding the Switch Inventory

Knowing your switch model and serial numbers can help TAC support with troubleshooting your fabric. This section explains how to find the switch model and serial numbers using the Cisco APIC GUI, CLI, and REST API.

Finding Your Switch Inventory Using the GUI

This section explains how to find your switch model and serial numbers using the Cisco APIC GUI.

Before You Begin

You must have access to the Cisco APIC GUI

Step 1 On the menu bar, choose **Fabric > Inventory**.

Step 2 In the navigation pane, click a **Pod** icon.
Your switch icons appear in the navigation pane.

Step 3 In the navigation pane, click on a switch icon.
A list of tabs appears at the top of the work pane.

Step 4 Click the **General** tab.
Your switch information appears in the work pane.

Finding Your Switch Inventory Using the NX-OS CLI

This section explains how to find your switch model and serial numbers using the NX-OS CLI.

Find your switch inventory as follows:

Example:

```
switch# show hardware
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Documents: http://www.cisco.com/en/US/products/ps9372/tsd_products_support_series_home.html
Copyright (c) 2002-2014, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php

Software
  BIOS:      version 07.56
  kickstart: version 12.1(1h) [build 12.1(1h)]
  system:    version 12.1(1h) [build 12.1(1h)]
  PE:        version 2.1(1h)
  BIOS compile time:       06/08/2016
  kickstart image file is: /bootflash/aci-n9000-dk9.12.1.1h.bin
  kickstart compile time:  10/01/2016 20:10:40 [10/01/2016 20:10:40]
  system image file is:    /bootflash/auto-s
  system compile time:     10/01/2016 20:10:40 [10/01/2016 20:10:40]


Hardware
  cisco N9K-C93180YC-EX ("supervisor")
   Intel(R) Xeon(R) CPU  @ 1.80GHz with 16400384 kB of memory.
  Processor Board ID FDO20101H1W

  Device name: ifav41-leaf204
  bootflash:    62522368 kB

Kernel uptime is 02 day(s), 21 hour(s), 42 minute(s), 31 second(s)

Last reset at 241000 usecs after Sun Oct 02 01:27:25 2016
  Reason: reset-by-installer
  System version: 12.1(1e)
  Service: Upgrade

plugin
  Core Plugin, Ethernet Plugin
--------------------------------
Switch hardware ID information
--------------------------------

Switch is booted up
Switch type is : Nexus C93180YC-EX Chassis
Model number is N9K-C93180YC-EX
H/W version is 0.2010
Part Number is 73-15298-01
Part Revision is 1
Manufacture Date is Year 20 Week 10
Serial number is FDO20101H1W
CLEI code is 73-15298-01

--------------------------------
```

```
Chassis has one slot
--------------------------------

Module1 ok
  Module type is : 48x10/25G
  1 submodules are present
  Model number is N9K-C93180YC-EX
  H/W version is 0.2110
  Part Number is 73-17776-02
  Part Revision is 11
  Manufacture Date is Year 20 Week 10
  Serial number is FDO20101H1W
  CLEI code is 73-17776-02

GEM ok
  Module type is : 6x40/100G Switch
  1 submodules are present
  Model number is N9K-C93180YC-EX
  H/W version is 0.2110
  Part Number is 73-17776-02
  Part Revision is 11
  Manufacture Date is Year 20 Week 10
  Serial number is FDO20101H1W
  CLEI code is 73-17776-02

---------------------------------------
Chassis has  2 PowerSupply Slots
---------------------------------------

PS1 shut
  Power supply type is : 54.000000W 220v AC
  Model number is NXA-PAC-650W-PE
  H/W version is 0.0
  Part Number is 341-0729-01
  Part Revision is A0
  Manufacture Date is Year 19 Week 50
  Serial number is LIT19500ZEK
  CLEI code is 341-0729-01

PS2 ok
  Power supply type is : 54.000000W 220v AC
  Model number is NXA-PAC-650W-PE
  H/W version is 0.0
  Part Number is 341-0729-01
  Part Revision is A0
  Manufacture Date is Year 19 Week 50
  Serial number is LIT19500ZEA
  CLEI code is 341-0729-01

---------------------------------------
Chassis has  4 Fans
---------------------------------------

FT1 ok

 Fan1(sys_fan1)(fan_model:NXA-FAN-30CFM-F)                                    is inserted but info
is not available

FT2 ok

 Fan2(sys_fan2)(fan_model:NXA-FAN-30CFM-F)                                    is inserted but info
is not available

FT3 ok

 Fan3(sys_fan3)(fan_model:NXA-FAN-30CFM-F)                                    is inserted but info
is not available

FT4 ok

 Fan4(sys_fan4)(fan_model:NXA-FAN-30CFM-F)                                    is inserted but info
```

```
is not available

==================================================================================
```

Finding Your Switch Inventory Using the REST API

This section explains how to find your switch model and serial numbers using the REST API

Find your switch inventory as follows:

Example:

```
GET
https://192.0.20.123/api/node/mo/topology/pod-1.json?query-target=children&target-subtree-class=fabricNode
```

The following response is returned:

```
response:
  {
     "totalCount":"8",
     "imdata":
     [{
         "fabricNode":{
           "attributes":{
              "adSt":"on",
              "childAction":"",
              "delayedHeartbeat":"no",
              "dn":"topology/pod-1/node-103",
              "fabricSt":"active",
              "id":"103",
              "lcOwn":"local",
              "modTs":"2016-10-08T14:49:35.665+00:00",
              "model":"N9K-C9396PX",
              "monPolDn":"uni/fabric/monfab-default",
              "name":"leaf3",
              "nameAlias":"",
              "role":"leaf",
              "serial":"TEP-1-103",
              "status":"","uid":"0",
              "vendor":"Cisco Systems, Inc",
              "version":""}
             }
     },{
         "fabricNode":{
           "attributes":{
             "adSt":"on",
             "childAction":"",
             "delayedHeartbeat":"no",
             "dn":"topology/pod-1/node-105",
             "fabricSt":"active",
        "id":"105",
        "lcOwn":"local",
        "modTs":"2016-10-08T14:47:52.011+00:00",
        "model":"N9K-C9508",
        "monPolDn":"uni/fabric/monfab-default",
        "name":"spine2",
        "nameAlias":"",
        "role":"spine",
        "serial":"TEP-1-105","status":"",
        "uid":"0",
        "vendor":"Cisco Systems, Inc",
        "version":""
     ...
  [TRUNCATED]
     ...
}
```

INDEX