Management

This chapter contains the following sections:

Management Workflows

ACI Management Access Workflows

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

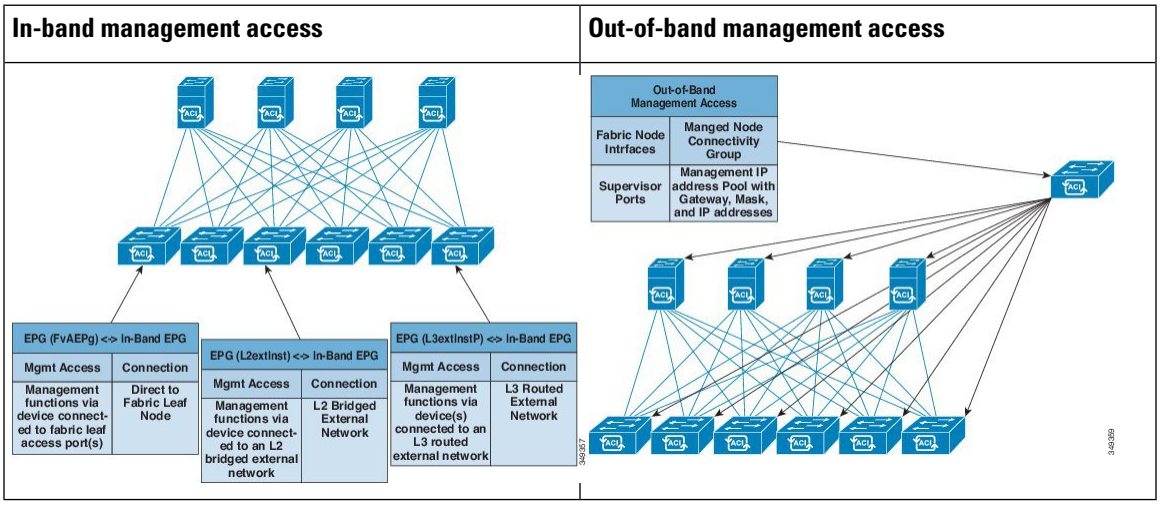

1. Prerequisites

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

2. Configure the ACI Leaf Switch Access Ports

Choose which of these management access scenarios you will use:

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

Suggested topics

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

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

Adding Management Access

Configuring the external management instance profile under the management tenant for in-band has no effect on the protocols that are configured under the fabric-wide communication policies. The subnets and contracts specified under the external management instance profile do not affect HTTP/HTTPS or SSH/Telnet.

Adding Management Access in the GUI

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

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

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

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

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

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

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

IPv4/IPv6 Addresses and In-Band Policies

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

IPv4/IPv6 Addresses in Out-of-Band Policies

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

IPv6 Table Modifications to Mirror the Existing IP Tables Functionality

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

Existing IP Tables

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

Modifications to IP Tables

1. When IP tables are created, they are first written into hash maps that are then written into intermediate file IP tables-new which are restored. When saved, a new IP tables file is created in the /etc/sysconfig/ folder. You can find both these files at the same location. Instead of making a system call for every rule, you must make a system call only while restoring and saving the file.

2. When a new policy is added instead of appending it to the file, an IP table is created from scratch, that is by loading default policies into the hashmaps, checking for new policies, and adding them to hashmaps. Later, they are written to the intermediate file (/etc/sysconfig/iptables-new) and saved.
3. It is not possible to configure source ports alone for a rule in out-of-band policy. Either destination port or source port along with a destination port can be added to the rules.
4. When a new policy is added, a new rule will be added to the IP tables file. This rule changes the access flow of IP tables default rules.

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

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

Note

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

Management Access Guidelines and Restrictions

- vzAny is supported as a consumer of a shared service but is not supported as a provider of a shared service. vzAny shared service consumer and vzAny provider is not supported.
- When configuring out-of-band management access, logging options for an out-of-band contract (enabling and viewing ACL contract and permit/deny logs) is not supported.
- An in-band management address must be configured for a leaf node in order to push the in-band management VRF to a leaf node.
- A bridge domain subnet IP address in the in-band management VRF can be assigned as a secondary IP address unless "Make this IP address primary" is selected for a gateway subnet.
- The following ports cannot be denied in an out-of-band contract:
  - 5010
  - 5012
  - 5013
  - 5020
  - 5021
  - 5025
  - 7777
  - 32768 through 60999

- A spine switch does not resolve ARP on the in-band mangement IP address. Due to this, any device in the in-band management network cannot communicate with the spine switch. Access to a spine switch is only possible over a Layer 3 network.
- With out-of-band management, the ICMP port is opened by default for all subnets.

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

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

Procedure

Step 1 To configure **In-Band Management Access**, perform the following steps:

a) On the menu bar, click **Tenants** > **mgmt**.
b) Expand **Quick Start**.
c) Click **In-Band Management Access** > **Configure In-Band Management Access** > **Start**.
d) Follow the instructions to add the **Nodes** in the management network, the **IP addresses** for the nodes, communication filters for the **Connected Devices**, and communication filters for **Remote Attached Devices**.

Step 2 To configure **Out-of-Band Management Access**, perform the following steps:

a) On the menu bar, click **Tenants** > **mgmt**.
b) Expand **Quick Start**.
c) Click **Out-of-Band Management Access** > **Configure Out-of-Band Management Access** > **Start**.
d) Follow the instructions to add the **Nodes** in the out-of-band management network, the **IP addresses** for the nodes, subnets allowed for the **External Hosts**, and communication filters that will determine communication for **Access**.

Configuring In-Band Management Access Using the Cisco APIC GUI

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

Procedure

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

Step 2 In the **Navigation** pane, right-click **Interfaces** and choose **Configure Interface, PC and VPC**.

Step 3 In the **Configure Interface, PC, and VPC** dialog box, to configure switch ports connected to Cisco Application Policy Infrastructure Controllers (APICs), perform the following actions:

a) Click the large + icon next to the switch diagram to create a new profile and configure VLANs for the Cisco APIC.
b) From the **Switches** field drop-down list, check the check boxes for the switches to which the Cisco APICs are connected. (leaf1 and leaf2).
c) In the **Switch Profile Name** field, enter a name for the profile (apicConnectedLeaves).
d) Click the + icon to configure the ports.
e) Verify that in the **Interface Type** area, the **Individual** radio button is selected.
f) In the **Interfaces** field, enter the ports to which the Cisco APICs are connected.
g) In the **Interface Selector Name** field, enter the name of the port profile (apicConnectedPorts).
h) In the **Interface Policy Group** field, click the **Create One** radio button.
i) In the **Attached Device Type** field, choose the appropriate device type to configure the domain (Bare Metal).
j) In the **Domain** field, click the **Create One** radio button.
k) In the **Domain Name** field, enter the domain name. (**inband**)
l) In the **VLAN** field, choose the **Create One** radio button.
m) In the **VLAN Range** field, enter the VLAN range. Click **Save**, and click **Save** again. Click **Submit**.

Step 4 In the **Navigation** pane, right-click **Switch Policies** and choose **Configure Interface, PC and VPC**.

Step 5 In the **Configure Interface, PC, and VPC** dialog box, perform the following actions:

a) Click the large + icon next to the switch diagram to create a new profile and configure VLANs for the server.
b) In the **Switches** field, from drop-down list, check the check boxes for the switches to which the servers are connected. (leaf1).
c) In the **Switch Profile Name** field, enter a name for the profile (vmmConnectedLeaves).
d) Click the + icon to configure the ports.
e) Verify that in the **Interface Type** area, the **Individual** radio button is selected.
f) In the **Interfaces** field, enter the ports to which the servers are connected (1/40).
g) In the **Interface Selector Name** field, enter the name of the port profile.
h) In the **Interface Policy Group** field, click the **Create One** radio button.
i) In the **Attached Device Type** field, choose the appropriate device type to configure the domain (Bare Metal).
j) In the **Domain** field, from the drop-down list click the **Choose One** radio button
k) From the **Physical Domain** drop-down list, choose the domain created earlier.
l) In the **Domain Name** field, enter the domain name.
m) Click **Save**, and click **Save** again.

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

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

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

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

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

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

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

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

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

a) In the **Create Node Management Addresses** dialog box, in the **Policy Name** field, enter the policy name (switchInb).
b) In the **Nodes** field, **Select** column, check the check boxes next to the nodes that will be part of this fabric (leaf1, leaf2, spine1, spine2).
c) In the **Config** field, click the **In-Band Addresses** checkbox.
d) In the **Node Range** fields, enter the range.
e) In the **In-Band IP Addresses** area, in the **In-Band Management EPG** field, from the drop-down list, choose default. The default in-band management EPG is now associated.
f) In the **In-Band IP Addresses** and **Gateway** fields, enter the IPv4 or IPv6 addresses as desired.
g) Click **Submit**. In the **Confirm** dialog box, click **Yes**. The IP addresses for the leaf and spine switches are now configured.

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

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

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

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

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

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

Before you begin

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

Procedure

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

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

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

a) In the **Policy Name** field, enter a policy name (switchOob).

b) In the **Nodes** field, check the check boxes next to the appropriate leaf and spine switches (leaf1, leaf2, spine1).

c) In the **Config** field, check the check box for **Out of-Band Addresses**.

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

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

e) In the **Out-Of-Band Gateway** field, enter the IP address and network mask for the external out-of-band management network.

f) In the **Out-of-Band IP Addresses** field, enter the range of desired IPv4 or IPv6 addresses that will be assigned to the switches. Click **Submit**.

The node management IP addresses are configured.

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

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

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

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

a) In the **Name** field, enter a name for the contract (oob-default).

b) Expand **Subjects**. In the **Create Contract Subject** dialog box, in the **Name** field, enter a subject name (oob-default).

c) Expand **Filter Chain**, and in the **Name** field, from the drop-down list, choose the name of the filter (default). Click **Update**, and click **OK**.

d) In the **Create Out-of-Band Contract** dialog box, click **Submit**.

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

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

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

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

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

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

a) In the **Name** field, enter a name (oob-mgmt-ext).

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

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

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

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

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

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

Exporting Tech Support, Statistics, and Core Files

About Exporting Files

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

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

File Export Guidelines and Restrictions

- HTTP export and the streaming API format is supported only with statistics information. Core and tech support data are not supported.
- The destination IP address for exported files cannot be an IPv6 address.
- Do not trigger tech support from more than five nodes simultaneously, especially if they are to be exported into the Cisco Application Policy Infrastructure Controller (APIC) or to an external server with insufficient bandwidth and compute resources.

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

Creating a Remote Location for Exporting Files

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

Procedure

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

Step 2 In the submenu bar, click **Import/Export**.

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

Step 4 Right-click **Remote Locations** and choose **Create Remote Path of a File**.

Step 5 In the **Create Remote Path of a File** dialog box, perform the following actions:

a) In the **Name** field, enter a name for the remote location.
b) In the **Host Name/IP** field, enter an IP address or a fully qualified domain name for the destination host.
c) In the **Protocol** field, click the radio button for the desired file transfer protocol.
d) In the **Remote Path** field, type the path where the file will be stored on the remote host.
e) Enter a username and password for logging in to the remote host and confirm the **Password**.
f) From the **Management EPG** drop-down list, choose the management EPG.
g) Click **Submit**.

Sending an On-Demand Tech Support File Using the GUI

Procedure

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

Step 2 In the submenu bar, click **Import/Export**.

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

Step 4 Right-click **On-demand Tech Support** and choose **Create On-demand Tech Support**.

The **Create On-demand Tech Support** dialog box appears.

Step 5 Enter the appropriate values in the fields of the **Create On-demand Tech Support** dialog box.

Note For an explanation of a field, click the help icon in the **Create On-demand Tech Support** dialog box. The help file opens to a properties description page.

Step 6 Click **Submit** to send the tech support file.

Note On-demand tech support files can be saved to another APIC to balance storage and CPU requirements. To verify the location, click on the On-demand Tech Support policy in the **Navigation** pane, then click the **OPERATIONAL** tab in the **Work** pane. The controller is displayed in the **EXPORT LOCATION** field.

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

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

Overview

This topic provides information on:

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

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

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

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

Configuration File Encryption

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

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

Observe the following guidelines when working with encrypted configuration files:

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

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

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

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

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

default setting is turned off could result in all the passwords of the ACI fabric to be removed upon the import of a configuration file that does not match the AES encryption settings of the fabric.

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

Configuring a Remote Location Using the GUI

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

Procedure

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

Step 2 In the navigation pane, right-click **Remote Locations** and choose **Create Remote Location**. The **Create Remote Location** dialog appears.

Step 3 Enter the appropriate values in the **Create Remote Location** dialog fields.

Note For an explanation of a field, click the 'i' icon to display the help file.

Step 4 When finished entering values in the **Create Remote Location** dialog fields, click **Submit**. You have now created a remote location for backing up your data.

Configuring an Export Policy Using the GUI

This procedure explains how to configure an Export policy using the APIC GUI. Follow these steps to trigger a backup of your data.

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

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

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

Procedure

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

Step 2 In the navigation pane, right-click **Export Policies** and choose **Create Configuration Export Policy**. The **Create Configuration Export Policy** dialog appears.

Step 3 Enter the appropriate values in the **Create Configuration Export Policy** dialog fields.

Note For an explanation of a field, click the help (**?**) icon to display the help file.

Step 4 When finished entering values in the **Create Configuration Export Policy** dialog fields, click **Submit**. You have now created a backup. You can view this under the **Configuration** tab (The backup file will appear in the **Configuration** pane on the right).

Note When deployed, and configured to do so, the Cisco Network Assurance Engine (NAE) creates export policies in the Cisco APIC for collecting data at timed intervals. You can identify an NAE export policy by its name, which is based on the assurance control configuration. If you delete an NAE export policy in the Cisco APIC, the NAE export policy will reappear in the Cisco APIC. We recommend not deleting the NAE export policies.

There's an **Operational** tab where you can see if it's running, successful, or failed. If you didn't trigger it yet, it is empty. If you created a backup, it creates a file that is shown in the **Operational** view of the backup file that was created. If you want to then import that data, you must create an Import policy.

Configuring an Import Policy Using the GUI

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

Procedure

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

Step 2 In the navigation pane, right-click **Import Policies** and click **Create Configuration Import Policy**. The **Create Configuration Import Policy** dialog appears.

Step 3 Enter the appropriate values in the **Create Configuration Import Policy** dialog fields.

Note For an explanation of a field, click the 'i' icon to display the help file. For more detailed information on import types and modes including (**Replace, Merge, Best Effort,** and **Atomic**), refer to the *Cisco Application Centric Infrastructure Fundamentals Guide* .

Step 4 When finished entering values in the **Create Configuration Import Policy** dialog fields, click **Submit**.

Note If you perform a clean reload of the fabric and import a previously-saved configuration, the time zone will change to UTC by default. Reset the time zone to your local time zone after the configuration import for the APIC cluster in these situations.

Encrypting Configuration Files Using the GUI

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

This section explains how to enable AES-256 encryption.

Procedure

Step 1 On the menu bar, choose **ADMIN > AAA**.

Step 2 In the navigation pane, click **AES Encryption Passphrase and Keys for Config Export (and Import)**. The **Global AES Encryption Settings for all Configurations Import and Export** window appears in the right pane.

Step 3 Create a passphrase, which can be between 16 and 32 characters long. There are no restrictions on the type of characters used.

Step 4 Click **SUBMIT**.

Note Once you have created and posted the passphrase, the keys are then generated in the back-end and the passphrase is not recoverable. Therefore, your passphrase is not visible to anyone because the key is automatically generated then deleted. Your backup only works if you know the passphrase (no one else can open it).

The **Key Configured** field now shows **yes**. You now see an encrypted hash (which is not the actual passphrase, but just a hash of it) in the **Encrpyted Passphrase** field.

Step 5 After setting and confirming your passphrase, check the check box next to **Enable Encryption** to turn the AES encryption feature on (checked).

The **Global AES Encryption Settings** field in your export and import policies will now be enabled by default.

Note

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

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

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

Configuration Import Behavior Scenario	Result
	fails to decrypt correctly, the import fails without importing any shards. Therefore, if a shard fails to decrypt all properties, all shards are rejected.

Backing up, Restoring, and Rolling Back Controller Configuration

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

Backing Up, Restoring, and Rolling Back Configuration Files Workflow

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

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

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

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

- Execution time
- Name of the file being processed/generated
- Status, as follows:
  - Pending
  - Running
  - Failed
  - Fail-no-data
  - Success
  - Success-with-warnings
- Details string (failure messages and warnings)

- Progress percentage = 100 * lastStepIndex/totalStepCount
- Field lastStepDescr indicating what was being done last

About the fileRemotePath Object

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

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

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

Sample Configuration

The following is a sample configuration:

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

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

Configuration Export to Controller

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

Note See the Snapshots section for more details.

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

- **name** - policy name
- **format** - format in which the data is stored inside the exported archive (xml or json)
- **targetDn** - the domain name (DN) of the specific object you want to export (empty means everything)
- **snapshot** - when true, the file is stored on the controller, no remote location configuration is needed

- **includeSecureFields** - Set to true by default, indicates whether the encrypted fields (passwords, etc.) should be included in the export archive.

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

Scheduling Exports

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

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

Troubleshooting

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

Sample Configuration Using the NX-OS Style CLI

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

```
apic1(config)# snapshot
 download  Configuration snapshot download setup mode
export    Configuration export setup mode
import    Configuration import setup mode
rollback  Configuration rollback setup mode
upload    Configuration snapshot upload setup mode
apic1(config)# snapshot export policy-name
apic1(config-export)#
 format    Snapshot format: xml or json
no        Negate a command or set its defaults
remote    Set the remote path configuration will get exported to
schedule  Schedule snapshot export
target    Snapshot target

bash      bash shell for unix commands
end       Exit to the exec mode
exit      Exit from current mode
fabric    show fabric related information
show      Show running system information
where     show the current mode
apic1(config-export)# format xml
apic1(config-export)# no remote path             [If no remote path is specified, the file
 is exported locally to a folder in the controller]
apic1(config-export)# target                     [Assigns the target of the export, which
can be fabric, infra, a specific tenant, or none. If no target is specified, all configuration
 information is exported.]
WORD  infra, fabric or tenant-x
apic1(config-export)#
apic1# trigger snapshot export policy-name       [Executes the snapshot export task]
apic1# ls /data2                                 [If no remote path is specified, the
configuration export file is saved locally to the controller under the folder data2]
ce_Dailybackup.tgz
```

Sample Configuration Using the GUI

The following is a sample configuration using the GUI:

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

Sample Configuration Using REST API

The following is a sample configuration using the REST API:

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

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

Configuration Import to Controller

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

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

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

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

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

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

- **importType**
  - replace - Current system configuration is replaced with the contents or the archive being imported (only atomic mode is supported)
  - merge - Nothing is deleted, archive content is applied on top the existing system configuration.
- **snapshot** - when true, the file is taken from the controller and no remote location configuration is needed.
- **failOnDecryptErrors** - (true by default) the file fails to import if the archive was encrypted with a different key than the one that is currently set up in the system.

Troubleshooting

The following scenarios may need troubleshooting:

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

Sample Configuration Using the NX-OS Style CLI

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

```
apic1# configure
apic1(config)# snapshot
 download  Configuration snapshot download setup mode
export    Configuration export setup mode
import    Configuration import setup mode
rollback  Configuration rollback setup mode
upload    Configuration snapshot upload setup mode
apic1(config)# snapshot import
 WORD       Import configuration name
default
rest-user
apic1(config)# snapshot import policy-name
apic1(config-import)#
 action  Snapshot import action merge|replace
file    Snapshot file name
mode    Snapshot import mode atomic|best-effort
no      Negate a command or set its defaults
remote  Set the remote path configuration will get imported from

bash    bash shell for unix commands
end     Exit to the exec mode
exit    Exit from current mode
fabric  show fabric related information
show    Show running system information
where   show the current mode
apic1(config-import)# file < from "show snapshot files" >
apic1(config-import)# no remote path
apic1(config-import)#
apic1# trigger snapshot import policy-name       [Executes the snapshot import task]
```

Sample Configuration Using the GUI

The following is a sample configuration using the GUI:

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

   of your configuration file.
10. Click **SUBMIT** to trigger your configuration import.

Sample Configuration Using the REST API

The following shows a sample configuration using the REST API:

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

Snapshots

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

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

configSnapshot objects provide the following:

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

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

Snapshot Manager Policy

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

You can also create a recurring snapshot.

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

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

Troubleshooting

For troubleshooting, refer to the Connectivity Issues section.

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

```
apic1(config)# snapshot upload policy-name
apic1(config-upload)#
 file    Snapshot file name
no      Negate a command or set its defaults
remote  Set the remote path configuration will get uploaded to

bash    bash shell for unix commands
end     Exit to the exec mode
exit    Exit from current mode
fabric  show fabric related information
show    Show running system information
where   show the current mode
apic1(config-upload)# file <file name from "show snapshot files">
apic1(config-upload)# remote path remote-path-name
apic1# trigger snapshot upload policy-name       [Executes the snapshot upload task]
```

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

```
apic1(config)# snapshot download policy-name
apic1(config-download)#
 file    Snapshot file name
no      Negate a command or set its defaults
remote  Set the remote path configuration will get downloaded from

bash    bash shell for unix commands
end     Exit to the exec mode
exit    Exit from current mode
fabric  show fabric related information
show    Show running system information
where   show the current mode
apic1(config-download)# file < file from remote path>
apic1(config-download)# remote path remote-path-name
apic1# trigger snapshot download policy-name       [Executes the snapshot download task]
```

Snapshot Upload and Download Using the GUI

To upload a snapshot file to a remote location:

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

To download a snapshot file from a remote location:

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

Snapshot Upload and Download Using the REST API

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

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

Rollback

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

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

Note

- The rollback feature only operates on snapshots.
- Remote archives are not supported directly. However, you can turn a remotely saved export into a snapshot using the snapshot manager policy (configSnapshotMgrP). For more information, see the Snapshot Manager Policy, on page 24
- The configRollbackP policy does not require a remote path configuration. If a remote path is provided, it will be ignored.

Rollback Workflow

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

The MOs are handled as follows:

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

The rollback feature also generates a diff file that contains the confiuration generated as a result of these calculations. Applying this configuration is the last step of the rollback process. The content of this file can be retrieved through a special REST API called readiff: apichost/mqapi2/snapshots.readiff.xml?jobdn=SNAPSHOT_JOB_DN.

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

Diff Tool

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

Sample Configuration Using the NX-OS Style CLI

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

```
apic1# show snapshot files
File    : ce2_DailyAutoBackup-2015-11-21T01-00-17.tar.gz
Created : 2015-11-21T01:00:21.167+00:00
Root    :
Size    : 22926

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

apic1# configure
apic1(config)# snapshot rollback myRollbackPolicy
apic1(config-rollback)# first-file ce2_DailyAutoBackup-2015-11-21T01-00-17.tar.gz
apic1(config-rollback)# second-file ce2_DailyAutoBackup-2015-11-21T09-00-21.tar.gz
apic1(config-rollback)# preview
apic1(config-rollback)# end
apic1# trigger snapshot rollback myRollbackPolicy
```

Sample Configuration Using the GUI

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

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

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

Sample Configuration Using the REST API

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

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

Using Syslog

About Syslog

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

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

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

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

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

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

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

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

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

Creating a Syslog Destination and Destination Group

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

Procedure

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

Step 2 In the submenu bar, click **External Data Collectors**.

Step 3 In the **Navigation** pane, expand **Monitoring Destinations**.

Step 4 Right-click **Syslog** and choose **Create Syslog Monitoring Destination Group**.

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

a) In the group and profile **Name** field, enter a name for the monitoring destination group and profile.

b) In the group and profile **Format** field, choose the format for Syslog messages.

The default is **aci**, or the RFC 3164 compliant message format, but you can choose to set it to the NX-OS style format instead.

c) In the group and profile **Admin State** drop-down list, choose **enabled**.

d) 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`.

e) 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.

f) Click **Next**.

g) In the **Create Remote Destinations** area, click **+** to add a remote destination.

Caution There is a risk of hostname resolution failure for remote syslog destinations if the DNS server that you specify is configured to be reachable over in-band connectivity. To avoid the issue, configure the syslog server using the IP address, or if you use a hostname, ensure that the DNS server is reachable over an out-of-band interface.

Step 6 In the **Create Syslog Remote Destination** dialog box, perform the following actions:

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 **Facility**.

The **Facility** is a number that you can optionally use to indicate which process generated the message, and can then be used to determine how the message will be handled at the receiving end.

e) In the 5.2(3) release and later, in the **Transport** field, choose the transport protocol to use for the messages.

- For releases prior to release 5.2(4), choose either **tcp** or **udp** as the transport protocol to use for the messages.
- In the 5.2(4) release and later, **ssl** is also an option for the transport protocol to use for the messages. This feature enables a Cisco ACI switch (acting as a client) to make a secure, encrypted outbound connection to remote syslog servers (acting as a server) supporting secure connectivity for logging. With authentication and encryption, this feature allows for a secure communication over an insecure network.

Note that you must also upload the necessary SSL certificate if you select **ssl** as the transport protocol to use for the messages. You can upload the necessary SSL certificate by navigating to the **Create Certificate Authority** window:

Admin > **AAA** > **Security** > **Public Key Management** > **Certificate Authorities**, then **Actions** > **Create Certificate Authority**

The default option for the transport protocol is **udp**.

f) From the **Management EPG** drop-down list, choose the management endpoint group.
g) Click **OK**.

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

Step 8 Click **Finish**.

Creating a Syslog Source

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

Before you begin

Create a syslog monitoring destination group.

Procedure

Step 1 From the menu bar and the navigation frame, navigate to a **Monitoring Policies** menu for the area of interest.

You can configure monitoring policies for tenants, fabric, and access.

Step 2 Expand **Monitoring Policies**, then select and expand a monitoring policy.

Under **Fabric > Fabric Policies > Monitoring Policies > Common Policy** is a basic monitoring policy that applies to all faults and events and is automatically deployed to all nodes and controllers in the fabric. Alternatively, you can specify an existing policy with a more limited scope.

Step 3 Under the monitoring policy, click **Callhome/SNMP/Syslog**.

Step 4 In the **Work** pane, choose **Syslog** from the **Source Type** drop-down list.

Step 5 From the **Monitoring Object** list, choose a managed object to be monitored.

If the desired object does not appear in the list, follow these steps:

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

Using Atomic Counters

About Atomic Counters

Atomic counters allow you to gather statistics about traffic between flows. Using atomic counters, you can detect drops and misrouting in the fabric, enabling quick debugging and isolation of application connectivity issues. For example, an administrator can enable atomic counters on all leaf switches to trace packets from endpoint 1 to endpoint 2. If any leaf switches have nonzero counters, other than the source and destination leaf switches, an administrator can drill down to those leafs.

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

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

- Counts of sent, received, dropped, and excess packets
  - Sent packets: The sent number reflects how many packets were sent from the source TEP (tunnel endpoint) to the destination TEP.
  - Received packets: The received number reflects how many packets the destination TEP received from the source TEP.
  - Dropped packets: The dropped number reflects how many packets were dropped during transmission. This number is the difference in the amount of packets sent and the amount of packets received.
  - Excess packets: The excess number reflects how many extra packets were received during transmission. This number is the amount of packets that were unexpectedly received due to a forwarding mismatch or a misrouting to the wrong place.

- Short-term data collection such as the last 30 seconds, and long-term data collection such as 5 minutes, 15 minutes, or more
- A breakdown of per-spine traffic (available when the number of TEPs, leaf or VPC, is less than 64)
- Ongoing monitoring

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

Tenant atomic counters can provide the following:

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

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

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

Atomic Counters Guidelines and Restrictions

- Use of atomic counters is not supported when the endpoints are in different tenants or in different contexts (VRFs) within the same tenant.
- In Cisco APIC release 3.1(2m) and later, if no statistics have been generated on a path in the lifetime of the fabric, no atomic counters are generated for the path. Also, the **Traffic Map** in the **Visualization** tab (**Operations** > **Visualization** in the Cisco APIC GUI) does not show all paths, only the active paths (paths that had traffic at some point in the fabric lifetime).
- In pure Layer 2 configurations where the IP address is not learned (the IP address is 0.0.0.0), endpoint-to-EPG and EPG-to-endpoint atomic counter policies are not supported. In these cases, endpoint-to-endpoint and EPG-to-EPG policies are supported. External policies are virtual routing and forwarding (VRF)-based, requiring learned IP addresses, and are supported.

- When the atomic counter source or destination is an endpoint, the endpoint must be dynamic and not static. Unlike a dynamic endpoint (fv:CEp), a static endpoint (fv:StCEp) does not have a child object (fv:RsCEpToPathEp) that is required by the atomic counter.
- In a transit topology, where leaf switches are not in full mesh with all spine switches, then leaf-to-leaf (TEP to TEP) counters do not work as expected.
- For leaf-to-leaf (TEP to TEP) atomic counters, once the number of tunnels increases the hardware limit, the system changes the mode from trail mode to path mode and the user is no longer presented with per-spine traffic.
- The atomic counter does not count spine proxy traffic.
- Packets dropped before entering the fabric or before being forwarded to a leaf port are ignored by atomic counters.
- Packets that are switched in the hypervisor (same Port Group and Host) are not counted.
- Atomic counters require an active fabric Network Time Protocol (NTP) policy.
- Atomic counters work for IPv6 sources and destinations, but configuring source and destination IP addresses across IPv4 and IPv6 addresses is not allowed.
- An atomic counter policy configured with fvCEp as the source or destination counts only the traffic that is from/to the MAC and IP addresses that are present in the fvCEp managed objects. If the fvCEp managed object has an empty IP address field, then all traffic to/from that MAC address would be counted regardless of the IP address. If the Cisco APIC has learned multiple IP addresses for an fvCEp, then traffic from only the one IP address in the fvCEp managed object itself is counted as previously stated. To configure an atomic counter policy to or from a specific IP address, use the fvIp managed object as the source or destination.
- If there is an fvIp behind an fvCEp, you must add fvIP-based policies and not fvCEp-based policies.
- Endpoint-to-endpoint atomic counter statistics are not reported for Layer 2 bridged traffic with IPv6 headers when the endpoints belong to the same EPG.
- For atomic counters to work for traffic flowing from an EPG or ESG to an L3Out EPG, configure the L3Out EPG with 0/1 and 128/1 to match all prefixes instead of 0/0.
- If your Cisco APIC has the traffic map mode set to "trial" and the Cisco APIC generated the F1545 fault, the only way that you can clear this fault is by setting the traffic map mode to "path." To change the traffic map mode, go to **Operations** > **Visualization**, click **Settings**, choose **path** for Mode, then click **Submit**. This will give you tunnel stats per port in both ingress and egress.

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

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

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

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

Configuring Atomic Counters

Procedure

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

Step 2 In the submenu bar, click the desired tenant.

Step 3 In the **Navigation** pane, expand the tenant and expand **Policies** and then expand **Troubleshoot**.

Step 4 Under **Troubleshoot**, expand **Atomic Counter Policy** and choose a traffic topology.

You can measure traffic between a combination of endpoints, endpoint groups, external interfaces, and IP addresses.

Step 5 Right-click the desired topology and choose **Add** *topology* **Policy** to open an **Add Policy** dialog box.

Step 6 In the **Add Policy** dialog box, perform the following actions:

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

b) choose or enter the identifying information for the traffic source.

The required identifying information differs depending on the type of source (endpoint, endpoint group, external interface, or IP address).

c) choose or enter the identifying information for the traffic destination.

d) (Optional) (Optional) In the **Filters** table, click the **+** icon to specify filtering of the traffic to be counted.

In the resulting **Create Atomic Counter Filter** dialog box, you can specify filtering by the IP protocol number (TCP=6, for example) and by source and destination IP port numbers.

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

Step 7 In the **Navigation** pane, under the selected topology, choose the new atomic counter policy.

The policy configuration is displayed in the **Work** pane.

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

Using SNMP

About SNMP

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

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

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

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

SNMP Access Support in Cisco ACI

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

SNMP support in Cisco ACI is as follows:

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

Note Cisco ACI supports a maximum of 10 trap receivers.

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

Table 1: SNMP Support Changes by Cisco APIC Release

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

SNMP Trap Aggregation

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

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

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

1. Configure each APIC controller to receive traps from the switches. Follow the procedure in Configuring an SNMP Trap Destination Using the GUI, on page 38 using the following settings:
   - In the **Host Name/IP** field, specify the IPv4 or IPv6 address of the APIC.
   - From the **Management EPG** list, select the out-of-band or inband management EPG.

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

2. Configure the APIC to forward aggregated traps to an external server. Follow the procedure in Configuring the SNMP Policy Using the GUI, on page 36 using the following settings:
   - In the **Trap Forward Servers** table, add the **IP Address** of the external server.

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

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

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

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

Configuring SNMP

Configuring the SNMP Policy Using the GUI

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

Before you begin

To allow SNMP communications, you must configure the following:

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

Procedure

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

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

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

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

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

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

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

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

This step is needed only if SNMPv3 access is required.

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

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

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

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

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

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

Step 8 Click **OK**.

Step 9 Click **Submit**.

Step 10 Under **Pod Policies**, expand **Policy Groups** and choose a policy group or right-click **Policy Groups** and choose **Create POD Policy Group**.

You can create a new pod policy group or you can use an existing group. The pod policy group can contain other pod policies in addition to the SNMP policy.

Step 11 In the pod policy group dialog box, perform the following actions:

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.

Procedure

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

Step 2 In the submenu bar, click **External Data Collectors**.

Step 3 In the **Navigation** pane, expand **Monitoring Destinations**.

Step 4 Right-click **SNMP** and choose **Create SNMP Monitoring Destination Group**.

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

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 IPv4 or IPv6 address or a fully qualified domain name for the destination host.
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 the **Security Name** and choose **noauth** as **v3 Security Level**.

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

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

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

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

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

Configuring an SNMP Trap Source Using the GUI

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

Procedure

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

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

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

You can create an SNMP source in the **Common Policy**, the **default** policy, or you can create a new monitoring policy.

Step 4 Expand the desired monitoring policy and choose **Callhome/SNMP/Syslog**.

If you chose the **Common Policy**, right-click **Common Policy**, choose **Create SNMP Source**, and follow the instructions below for that dialog box.

Step 5 In the **Work** pane, from the **Monitoring Object** drop-down list, choose **ALL**.

Step 6 From the **Source Type** drop-down list, choose **SNMP**.

Step 7 In the table, click the + icon to open the **Create SNMP Source** dialog box.

Step 8 In the **Create SNMP Source** dialog box, perform the following actions:

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.

Using SPAN

About SPAN

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

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

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

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

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

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

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

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

- You no longer have to include the IP prefix of the Layer 3 interface when configuring source SPAN with Layer 3 interface filtering.
- Support for configuring filter groups, which is a grouping of one or more filter entries. Use the filter group to specify the matching criteria that will be used to determine if a received packet should be analyzed using SPAN.
- The SPAN-on-drop feature, which captures packets that are dropped due to forwarding at the ingress in the ASIC and sends them to a pre-configured SPAN destination. There are 3 types of SPAN-on-drop configuration: access drop using access ports as a SPAN source, fabric drop using fabric ports as a SPAN source, and global drop using all ports on a node as a SPAN source. SPAN-on-drop is configured using regular SPAN (through the CLI, GUI, and REST API) and using troubleshooting SPAN (CLI and REST API, only). For more information about configuring this feature, see *Configuring SPAN Using the GUI*, *Configuring SPAN Using the NX-OS Style CLI*, and *Configuring SPAN Using the REST API*.

SPAN Guidelines and Restrictions

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

- Generation 1 switches are identified by the lack of a suffix, such as "EX", "FX", or "FX2," at the end of the switch name (for example, N9K-9312TX).
- Generation 2 switches are identified with a suffix, such as "EX", "FX", or "FX2," at the end of the switch name.

- The type of SPAN supported varies:
  - For generation 1 switches, tenant and access SPAN use the encapsulated remote extension of SPAN (ERSPAN) type I (Version 1 option in the Cisco Application Policy Infrastructure Controller (APIC) GUI).
  - For generation 2 switches, tenant and access SPAN use the encapsulated remote extension of SPAN (ERSPAN) type II (Version 2 option in the Cisco APIC GUI).
  - Fabric SPAN uses ERSPAN type II.
- When configuring ERSPAN session, if the SPAN source contains a destination and interfaces from a spine switch within a GOLF VRF instance, an L3Out prefix is sent to the GOLF router with the wrong BGP next-hop, breaking connectivity from GOLF to that L3Out.
- A uSeg EPG or ESG cannot be used as a SPAN source EPG because the SPAN source filter is based on the VLAN ID. Thus, even if an endpoint is classified to a uSeg EPG or an ESG, traffic from the endpoint is mirrored if its VLAN is the VLAN of the SPAN source EPG.
- You cannot specify an l3extLIfP Layer 3 subinterface as a SPAN source. You must use the entire port for monitoring traffic from external sources.
- In local SPAN for FEX interfaces, the FEX interfaces can only be used as SPAN sources, not SPAN destinations.
  - On generation 1 switches, Tx SPAN does not work for any Layer 3 switched traffic.
  - On generation 2 switches, Tx SPAN does not work whether traffic is Layer 2 or Layer 3 switched.

  There are no limitations for Rx SPAN.
- For SPAN of FEX fabric port-channel (NIF), the member interfaces are supported as SPAN source interfaces on generation 1 leaf switches.

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

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

- ERSPAN destination IP addresses must be learned in the fabric as an endpoint.
- SPAN supports IPv6 traffic but the destination IP address for the ERSPAN cannot be an IPv6 address.
- The individual port member of a port channel or a vPC cannot be configured as the source. Use the port channel, vPC, or vPC component as the source in the SPAN session.
- A fault is not raised on the ERSPAN source group when the destination EPG is deleted or unavailable.
- SPAN filters are supported on generation 2 leaf switches only.
- An access SPAN source supports only one of the following filters at a given time:
  - EPG
  - Routed outside (L3Out)
- When deploying the access SPAN source with an L3Out filter, ensure that the L3Out is also deployed on the matching interface:
  - If an L3Out is deployed on a port, a SPAN source must be deployed on the same port.
  - If an L3Out is deployed on a PC, a SPAN source must be deployed on the same PC.
  - If an L3Out is deployed on a vPC, a SPAN source must be deployed on the same vPC.
- An L3Out routed interface and routed sub-interface can be deployed on a port or a PC, but an L3Out SVI can be deployed on a port, PC, or vPC. A SPAN source with an L3Out filter must be deployed accordingly.
- An L3Out filter is not supported in fabric SPAN or tenant SPAN sessions.
- The correct L3Out must be selected in the L3 configuration tab of the EPG bridge domain; otherwise, packet flow for basic L3Out will not work.
- An encapsulation value is mandatory for a routed sub-interface and SVI, but is not applicable for a routed interface. The L3Out sub-interface or SVI encapsulation value must be different from the EPG encapsulation value.
- When an EPG filter is enabled within a SPAN session, ARP packets, which are sent out of the interface in the transit, or tx, direction, will not be spanned.
- SPAN filters are not supported in the following:
  - Fabric ports
  - Fabric and tenant SPAN sessions
  - Spine switches
- L4 port range filter entries will not be added if you attempt to add more L4 port ranges than are officially supported.
- A SPAN session will not come up if you attempt to associate more than the supported filter entries at the SPAN source group level or at the individual SPAN source level.
- Deleted filter entries will remain in TCAM if you add or delete more filters entries than are officially supported.

- See the *Verified Scalability Guide for Cisco ACI* document for SPAN-related limits, such as the maximum number of active SPAN sessions and SPAN filter limitations.
- For the SPAN-on-drop feature, the following guidelines and restrictions apply:
  - The SPAN-on-drop feature is supported on generation 2 leaf switches.
  - The SPAN-on-drop feature only captures packets with forwarding drops in the LUX block, which captures forwarding drop packets at the ingress. The SPAN-on-drop feature cannot capture the BMX (buffer) and RWX (egress) drops.
  - When using the troubleshooting CLI to create a SPAN session with SPAN-on-drop enabled and Cisco APIC as the destination, the session is disabled when 100 MB of data is captured.
  - On a modular chassis, the SPAN-on-drop feature will only work for the packets dropped on the line cards. Packets that are dropped on the fabric card will not be spanned.
  - SPAN-on-drop ACLs with other SPAN ACLs are not merged. If a SPAN-on-drop session is configured on an interface along with ACL-based SPAN, then any packets dropped on that interface will only be sent to the SPAN-on-drop session.
  - You cannot confgure SPAN on drop and SPAN ACL on the same session.
  - When an access or fabric port-drop session and a global-drop session are configured, the access or fabric port-drop session takes the priority over the global-drop session.
  - The number of filter entries supported in TCAM = (M * S1 * 1 + N * S2 * 2) + (S3 * 2). This is applicable to rx SPAN or tx SPAN, separately. Currently, the maximum filter entries supported in tx or rx SPAN is 480 in each direction when following this formula (and assuming there are no other sources that are configured without filter-group association [means S3 = 0] and with 16 port-ranges included). When the number of filter entries exceed the maximum number allowed, a fault will be raised. Note that you can specify Layer 4 port ranges in the filter entry. However, sixteen Layer 4 ports are programmed into the hardware as a single filter entry.

Note

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

- With MAC pinning configured in the LACP policy for a PC or vPC, the PC member ports will be placed in the LACP individual port mode and the PC is operationally non-existent. Hence, a SPAN source configuration with such a PC will fail, resulting in the generation of the "No operational src/dst" fault. With the MAC pinning mode configured, SPAN can be configured only on individual ports.
- A packet that is received on a Cisco Application Centric Infrastructure (ACI) leaf switch will be spanned only once, even if span sessions are configured on both the ingress and egress interfaces.
- When you use a routed outside SPAN source filter, you see only unicast in the Tx direction. In the Rx direction, you can see unicast, broadcast, and multicast.

- An L3Out filter is not supported for transmit multicast SPAN. An L3Out is represented as a combination of sclass/dclass in the ingress ACL filters and can therefore match unicast traffic only. Transmit multicast traffic can be spanned only on ports and port-channels.
- You can use a port channel interface as a SPAN destination only on -EX and later switches.
- You cannot configure multiple SPAN sessions with the same source interface when a SPAN filter (5-tuple filter) is applied.
- The local SPAN destination port of a leaf switch does not expect incoming traffic. You can ensure that the switch drops incoming SPAN destination port traffic by configuring a Layer 2 interface policy and setting the **VLAN Scope** property to **Port Local scope** instead of **Global scope**. Apply this policy to the SPAN destination ports. You can configure an Layer 2 interface policy by going to the following location in the GUI: **Fabric** > **Access Policies** > **Policies** > **Interface** > **L2 Interface**.
- When you configure SPAN for a given packet, SPAN is supported for the packet only once. If traffic is selected by SPAN in Rx for the first SSN, the traffic will not selected by SPAN again in Tx for a second SSN. Thus, when the SPAN session ingress and egress port sits on a single switch, the SPAN session capture will be one-way only. The SPAN session cannot display two-way traffic.
- A SPAN ACL filter configured in the filter group does not filter the broadcast, unknown-unicast and multicast (BUM) traffic that egresses the access interface. A SPAN ACL in the egress direction works only for unicast IPv4 or IPv6 traffic.
- When configuring a SPAN destination as a local port, EPGs cannot be deployed to that interface.
- In a leaf switch, a SPAN source with a VRF filter will match all regular bridge domains and all Layer 3 SVIs under the VRF instance.
- In a spine switch, a SPAN source with a VRF matches only the configured VRF VNID traffic and a bridge domain filter will match only the bridge domain VNID traffic.

Configuring a Tenant SPAN Session Using the Cisco APIC GUI

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

Procedure

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

Step 2 In the submenu bar, click the tenant that contains the source endpoint.

Step 3 In the **Navigation** pane, expand the tenant, expand **Policies** > **Troubleshooting** > **SPAN**.

Two nodes appear under **SPAN**: **SPAN Destination Groups** and **SPAN Source Groups**.

Step 4 From the **Navigation** pane, right-click **SPAN Source Groups** and choose **Create SPAN Source Group**. The **Create SPAN Source Group** dialog appears.

Step 5 Enter the appropriate values in the required fields of the **Create SPAN Source Group** dialog box.

Note For a description of a field, click the information icon (**?**) at the top-right corner of the dialog box to display the help file.

Step 6 Expand the **Create Sources** table to open the **Create SPAN Source** dialog box.

Step 7 Enter the appropriate values in the **Create SPAN Source** dialog box fields.

Note For the explanation of a field, click the help icon (**?**) to view the help file.

Step 8 When finished creating the SPAN source, click **OK**.

You return to the **Create SPAN Source Group** dialog box.

Step 9 When finished entering values in the **Create SPAN Source Group** dialog box fields, click **Submit**.

What to do next

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

Using Traceroute

About Traceroute

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

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

Traceroute supports a variety of modes, including:

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

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

Traceroute Guidelines and Restrictions

- When the traceroute source or destination is an endpoint, the endpoint must be dynamic and not static. Unlike a dynamic endpoint (fv:CEp), a static endpoint (fv:StCEp) does not have a child object (fv:RsCEpToPathEp) that is required for traceroute.

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

Performing a Traceroute Between Endpoints

Procedure

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

Step 2 In the submenu bar, click the tenant that contains the source endpoint.

Step 3 In the **Navigation** pane, expand the tenant and expand **Policies** > **Troubleshoot**.

Step 4 Under **Troubleshoot**, right-click on one of the following traceroute policies:

- **Endpoint-to-Endpoint Traceroute Policies** and choose **Create Endpoint-to-Endpoint Traceroute Policy**
- **Endpoint-to-External-IP Traceroute Policies** and choose **Create Endpoint-to-External-IP Traceroute Policy**
- **External-IP-to-Endpoint Traceroute Policies** and choose **Create External-IP-to-Endpoint Traceroute Policy**
- **External-IP-to-External-IP Traceroute Policies** and choose **Create External-IP-to-External-IP Traceroute Policy**

Step 5 Enter the appropriate values in the dialog box fields and click **Submit**.

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

Step 6 In the **Navigation** pane or the **Traceroute Policies** table, click the traceroute policy.

The traceroute policy is displayed in the **Work** pane.

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

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

Note

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