



Cisco APIC Basic Configuration Guide, Release 5.3(x)

First Published: 2023-11-24

Last Modified: 2024-01-19

Americas Headquarters

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



Trademarks

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at:

<http://www.cisco.com/go/softwareterms>. Cisco product warranty information is available at <http://www.cisco.com/go/warranty>. US Federal Communications Commission Notices are found here <http://www.cisco.com/c/en/us/products/us-fcc-notice.html>.

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

Any products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and-if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

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

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com go trademarks](http://www.cisco.com/go/trademarks). Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)



CONTENTS

PREFACE	Trademarks iii
----------------	-----------------------

CHAPTER 1	New and Changed Information 1
	New and Changed Information 1

CHAPTER 2	About Cisco ACI/APIC Configuration 3
	Recommended Settings for the Cisco Application Policy Infrastructure Controller 3
	About ACI/APIC Interfaces 5
	Mixing the NX-OS Style CLI and the APIC GUI 6
	About the Modes of Configuring Layer 3 External Connectivity 7
	Configuration Validation 8

CHAPTER 3	First Time Setup Wizard 9
	About the First Time Setup Wizard 9
	Fabric Membership 9
	Management 11
	vPC Pairs 11
	BGP 12
	DNS 12
	NTP 13
	Proxy 13
	Global Configurations 13

CHAPTER 4	User Access, Authentication, and Accounting 17
	Access Rights Workflow Dependencies 17
	User Access, Authorization, and Accounting 17

- Multiple Tenant Support 18
- User Access: Roles, Privileges, and Security Domains 18
- Login Domains 20
- Creating Login Domain Using the GUI 20
- Creating a Provider 21
- Configuring a Local User 25
 - Configuring a Local User Using the GUI 25
 - Configuring SSH Public Key Authentication Using the GUI 27
- Configuring a Remote User 27
 - AV Pair on the External Authentication Server 28
 - Configuring an AV Pair on the External Authentication Server 29
 - Configuring APIC for TACACS+ Access 30
 - Configuring APIC for RADIUS Access 31
 - Configuring a Cisco Secure Access Control Server for RADIUS and TACACS+ Access to the APIC 32
- Configuring Windows Server 2012 LDAP for APIC Access with Cisco AVPair 33
- Configuring APIC for LDAP Access 35
- Changing the Default Behavior for Remote Users with Missing or Bad Cisco AV Pairs 36
- About Signature-Based Transactions 36
 - Guidelines and Limitations 37
 - Generating an X.509 Certificate and a Private Key 37
 - Configuring a Local User 38
 - Creating a Local User and Adding a User Certificate Using the GUI 38
 - Creating a Local User Using Python SDK 39
 - Using a Private Key to Calculate a Signature 40
- Accounting 43
- Routed Connectivity to External Networks as a Shared Service Billing and Statistics 44

CHAPTER 5

Management 45

- Management Workflows 45
 - ACI Management Access Workflows 45
- Adding Management Access 47
 - Adding Management Access in the GUI 47
 - IPv4/IPv6 Addresses and In-Band Policies 47

IPv4/IPv6 Addresses in Out-of-Band Policies	47
IPv6 Table Modifications to Mirror the Existing IP Tables Functionality	48
Management Access Guidelines and Restrictions	48
Configuring In-Band and Out-of-Band Management Access with Wizards	49
Configuring In-Band Management Access Using the Cisco APIC GUI	50
Configuring Out-of-Band Management Access Using the Cisco APIC GUI	52
Exporting Tech Support, Statistics, and Core Files	54
About Exporting Files	54
File Export Guidelines and Restrictions	54
Creating a Remote Location for Exporting Files	55
Sending an On-Demand Tech Support File Using the GUI	55
Overview	56
Configuration File Encryption	56
Configuring a Remote Location Using the GUI	58
Configuring an Export Policy Using the GUI	58
Configuring an Import Policy Using the GUI	59
Encrypting Configuration Files Using the GUI	60
Backing up, Restoring, and Rolling Back Controller Configuration	62
Backing Up, Restoring, and Rolling Back Configuration Files Workflow	62
About the fileRemotePath Object	62
Configuration Export to Controller	63
Configuration Import to Controller	65
Snapshots	68
Snapshot Manager Policy	68
Rollback	70
Using Syslog	72
About Syslog	72
Creating a Syslog Destination and Destination Group	72
Creating a Syslog Source	74
Using Atomic Counters	75
About Atomic Counters	75
Atomic Counters Guidelines and Restrictions	76
Configuring Atomic Counters	78
Using SNMP	78

About SNMP	78
SNMP Access Support in Cisco ACI	79
SNMP Trap Aggregation	79
Configuring SNMP	80
Configuring the SNMP Policy Using the GUI	80
Configuring an SNMP Trap Destination Using the GUI	82
Configuring an SNMP Trap Source Using the GUI	83
Monitoring the System Using SNMP	83
Using SPAN	84
About SPAN	84
Multinode SPAN	85
SPAN Guidelines and Restrictions	85
Configuring SPAN Using the GUI	89
Configuring a Tenant SPAN Session Using the Cisco APIC GUI	89
Configuring a SPAN Filter Group Using the APIC GUI	90
Configuring an Access SPAN Policy Using the Cisco APIC GUI	91
Configuring a Fabric SPAN Policy Using the Cisco APIC GUI	92
Configuring a Layer 3 EPG SPAN Session for External Access Using the APIC GUI	92
Configuring a Destination Group for an Access SPAN Policy Using the Cisco APIC GUI	94
Configuring a Destination Group for a Fabric SPAN Policy Using the Cisco APIC GUI	94
Configuring SPAN Using the NX-OS-Style CLI	95
Configuring Local SPAN in Access Mode Using the NX-OS-Style CLI	95
Configuring a SPAN Filter Group Using the NX-OS-Style CLI	97
Associating a SPAN Filter Group Using the NX-OS-Style CLI	99
Configuring ERSPAN in Access Mode Using the NX-OS-Style CLI	100
Configuring ERSPAN in Fabric Mode Using the NX-OS-Style CLI	103
Configuring ERSPAN in Tenant Mode Using the NX-OS-Style CLI	106
Configuring a Global SPAN-On-Drop Session Using the NX-OS-Style CLI	108
Configuring SPAN Using the REST API	109
Configuring a Fabric Destination Group for an ERSPAN Destination Using the REST API	109
Configuring a Global Drop Source Group Using the REST API	110
Configuring a Leaf Port as a SPAN Destination Using the REST API	110
Configuring a SPAN Access Source Group Using the REST API	111
Configuring a SPAN Fabric Source Group Using the REST API	111

Configuring an Access Destination Group for an ERSPAN Destination Using the REST API	112
Using Traceroute	112
About Traceroute	112
Traceroute Guidelines and Restrictions	113
Performing a Traceroute Between Endpoints	113
acidiag Command	114

CHAPTER 6
Provisioning Core ACI Fabric Services 123

Link Level Policies	123
Electromagnetic Interference Retrain	123
Configuring a Link Level Policy Using the GUI	123
Port Bring-up Delay	124
Link Debounce Interval	124
Link Flap Policies	124
Configuring a Link Flap Policy Using the GUI	125
Time Synchronization and NTP	125
In-Band Management NTP	126
NTP over IPv6	126
Configuring NTP Using the GUI	126
Configuring NTP Using the REST API	127
Verifying NTP Operation Using the GUI	128
NTP Server	128
Enabling the NTP Server Using the GUI	129
Configuring the Datetime Format Using the GUI	130
Configuring a DHCP Relay Policy	131
Guidelines and Limitations for a DHCP Relay Policy	134
Configuring a DHCP Server Policy for the APIC Infrastructure Using the GUI	134
Configuring Option 79 Using REST API	137
Configuring a DHCP Server Policy for the APIC Infrastructure Using the NX-OS Style CLI	137
Configuring a DHCP Server Policy for the APIC Infrastructure Using the REST API	138
Configuring a DNS Service Policy	139
Configuring External Destinations with an In-Band DNS Service Policy	140
Dual Stack IPv4 and IPv6 DNS Servers	141
Dual-Stack IPv4 and IPv6 Environment	141

Policy for Priority of IPv4 or IPv6 in a DNS Profile	141
Configuring a DNS Service Policy to Connect with DNS Providers Using the GUI	142
Configuring Custom Certificates	143
Configuring Custom Certificate Guidelines	143
Configuring a Custom Certificate for Cisco ACI HTTPS Access Using the GUI	144
Provisioning Fabric Wide System Settings	146
Configuring APIC In-Band or Out-of-Band Connectivity Preferences	146
Configure Quota Management Policies	147
Create an Enforced BD Exception List	147
Create a BGP Route Reflector Policy and Route Reflector Node Endpoints	148
Configure a Fabric Wide Control Plane MTU Policy	149
Configure Endpoint Loop Protection	149
Rogue Endpoint Control Policy	150
About the Rogue Endpoint Control Policy	150
Limitations of the Rogue Endpoint Control Policy	150
Configuring the Rogue Endpoint Control Policy Using the GUI	151
Configure the Rogue Endpoint Control Policy Using the NX-OS Style CLI	152
Configure the Rogue Endpoint Control Policy Using the REST API	153
About the Rogue/COOP Exception List	153
Guidelines and Limitations for the Rogue/COOP Exception List	154
Configuring the Rogue/COOP Exception List While Creating a Bridge Domain Using the GUI	154
Configuring the Rogue/COOP Exception List of an Existing Bridge Domain Using the GUI	155
Configuring the Rogue/COOP Exception List of an Existing Bridge Domain Using the REST API	156
About Max IP Address Flow Control	156
Configure COOP	156
About COOP	156
Viewing COOP Dampened Endpoints Using the GUI	158
Viewing COOP Dampened Endpoints Using the Switch CLI	158
Clearing COOP Dampened Endpoints Using the GUI	159
Clearing a COOP Dampened Endpoint Using the Switch CLI	159
Disabling COOP Endpoint Dampening Using the REST API	159
Configuring COOP Authentication Using the APIC GUI	160
Configuring COOP Authentication Using the Cisco NX-OS-Style CLI	160

Configuring COOP Authentication Using the REST API	160
Endpoint Listen Policy	161
About the Endpoint Listen Policy	161
Configuring the Endpoint Listen Policy Using the GUI	161
Configure IP Aging	161
Disable Remote Endpoint Learning	162
Globally Enforce Subnet Checks	163
Reallocate a GIPo	163
Globally Enforce Domain Validation	164
Enable OpFlex Client Authentication	164
Fabric Load Balancing	165
Creating a Load Balancer Policy Using the Cisco APIC GUI	167
Creating a Load Balancer Policy Using the CLI	167
Creating a Load Balancer Policy Using the REST API	168
Enable a Time Precision Policy	169
Enable a Global System GIPo Policy	170
Configure a Fabric Port Tracking Policy	170
Provisioning Global Fabric Access Policies	171
Create a Global Attachable Access Entity Profile	171
Configure the Global QoS Class Policy	172
Create a Global DHCP Relay Policy	172
Enable a Global MCP Instance Policy	174
Create an Error Disabled Recovery Policy	174
Per Port Policies	175
About Per Port Policies	175
Configuring a Per Port Policy Using the GUI	175
Validating a Per Port Policy Using the GUI	176
Showing the Hidden Policies Using the GUI	177
Creating a Mis-cabling Protocol Interface Policy Using the GUI (Optional)	177

CHAPTER 7
Basic User Tenant Configuration 179

Tenants	179
Routing Within the Tenant	180
Layer 3 VNIDs Facilitate Transporting Inter-subnet Tenant Traffic	181

Router Peering and Route Distribution	183
Bridged Interface to an External Router	183
Configuring Route Reflectors	184
Configuring External Connectivity Using a Layer 3 Out	185
Configuring an MP-BGP Route Reflector Using the GUI	185
Configuring an MP-BGP Route Reflector for the ACI Fabric	185
Configuring an MP-BGP Route Reflector Using the REST API	186
Verifying the MP-BGP Route Reflector Configuration	187
Creating an OSPF L3Out for Management Tenant Using the GUI	187
Creating an OSPF External Routed Network for a Tenant Using the NX-OS CLI	189
Creating Tenants, VRFs, and Bridge Domains	191
Tenants Overview	191
Tenant Creation	191
VRF and Bridge Domains	191
Creating a Tenant, VRF, and Bridge Domain Using the GUI	192
Deploying EPGs	193
Statically Deploying an EPG on a Specific Port	193
Deploying an EPG on a Specific Node or Port Using the GUI	193
Creating Domains, Attach Entity Profiles, and VLANs to Deploy an EPG on a Specific Port	195
Creating Domains, and VLANs to Deploy an EPG on a Specific Port Using the GUI	195
Deploying an Application EPG through an AEP or Interface Policy Group to Multiple Ports	196
Deploying an EPG through an AEP to Multiple Interfaces Using the APIC GUI	197
Microsegmented EPGs	198
Using Microsegmentation with Network-based Attributes on Bare Metal	198
Configuring Network-based Microsegmented EPGs in a Bare-Metal environment Using the GUI	198
IP Address-Based Microsegmented EPG as a Shared Resource	200
Configuring an IP-based Microsegmented EPG as a Shared Resource Using the GUI	200
Unconfiguring an IP-based Microsegmented EPG as a Shared Resource Using the GUI	201
Deploying Application Profiles and Contracts	202
Security Policy Enforcement	202
Contracts Contain Security Policy Specifications	203
Three-Tier Application Deployment	205
Parameters to Create a Filter for http	206

Parameters to Create Filters for rmi and sql	207
Example Application Profile Database	207
Creating an Application Profile Using the GUI	207
Creating EPGs Using the GUI	208
Configuring Contracts Using the APIC GUI	209
Guidelines and Limitations for Contracts and Filters	209
Creating a Filter Using the GUI	209
Creating a Contract Using the GUI	210
Consuming and Providing Contracts Using the GUI	210
Optimize Contract Performance	211
Optimize Contract Performance	211
Configure a Contract with Optimized TCAM Usage Using the GUI	213
Contract and Subject Exceptions	214
Configuring Contract or Subject Exceptions for Contracts	214
Configure a Contract or Subject Exception Using the GUI	215
Intra-EPG Contracts	216
Intra-EPG Contracts	216
Guidelines and Limitations for Intra-EPG Contracts	216
Adding an Intra-EPG Contract to an Application EPG Using the GUI	217
Adding an Intra-EPG Contract to an Application EPG Using the NX-OS-Style CLI	219
Adding an Intra-EPG Contract to an Application EPG Using the REST API	220
EPG Contract Inheritance	224
About Contract Inheritance	224
Configuring EPG Contract Inheritance Using the GUI	226
Configuring Application EPG Contract Inheritance Using the GUI	226
Configuring uSeg EPG Contract Inheritance Using the GUI	226
Configuring L2Out EPG Contract Inheritance Using the GUI	227
Configuring External L3Out EPG Contract Inheritance Using the GUI	228
Contract Preferred Groups	228
About Contract Preferred Groups	228
Guidelines for Contract Preferred Groups	230
Configuring Contract Preferred Groups Using the GUI	230
Creating an L4-L7 Service EPG Policy Using the GUI	231
Contracts with Permit and Deny Rules	232

About Contracts with Permit and Deny Rules	232
Enabling ACL Contract Permit and Deny Logging Using the GUI	232
Enabling ACL Contract Permit Logging Using the NX-OS CLI	233
Enabling ACL Contract Permit Logging Using the REST API	234
Enabling Taboo Contract Deny Logging Using the GUI	234
Enabling Taboo Contract Deny Logging Using the NX-OS CLI	235
Enabling Taboo Contract Deny Logging Using the REST API	236
Viewing ACL Permit and Deny Logs Using the GUI	236
Viewing ACL Permit and Deny Logs Using the REST API	237
Viewing ACL Permit and Deny Logs Using the NX-OS CLI	238

APPENDIX A

Configuring the Cisco APIC Using the CLI	241
Cluster Management Guidelines	241
Replacing a Cisco APIC in a Cluster Using the CLI	242
Reducing the APIC Cluster Size	244
Contracting the Cisco APIC Cluster	244
Switching Over Active APIC with Standby APIC Using CLI	245
Verifying Cold Standby Status Using the CLI	245
Registering an Unregistered Switch Using the CLI	246
Adding a Switch Before Discovery Using the CLI	246
Removing a Switch to Maintenance Mode Using the CLI	247
Inserting a Switch to Operation Mode Using the CLI	247
Configuring a Remote Location Using the NX-OS Style CLI	247
Finding Your Switch Inventory Using the NX-OS CLI	248
Verifying the Cisco APIC Cluster Using the CLI	251

APPENDIX B

Configuring the Cisco APIC Using the REST API	255
Expanding the APIC Cluster Using the REST API	255
Contracting the APIC Cluster Using the REST API	256
Reducing the APIC Cluster Size	258
Switching Over Active APIC with Standby APIC Using REST API	258
Registering an Unregistered Switch Using the REST API	259
Adding a Switch Before Discovery Using the REST API	260
Removing a Switch to Maintenance Mode Using the REST API	260

Inserting a Switch to Operational Mode Using the REST API	261
Configuring a Remote Location Using the REST API	261
Sending an On-Demand Tech Support File Using the REST API	262
Finding Your Switch Inventory Using the REST API	262



CHAPTER 1

New and Changed Information

- [New and Changed Information](#), on page 1

New and Changed Information

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

Table 1: New Features and Changed Information for Cisco APIC Release 5.3(1)

Feature	Description	Where Documented
N/A	This document has no changes from the previous release.	N/A



CHAPTER 2

About Cisco ACI/APIC Configuration

- [Recommended Settings for the Cisco Application Policy Infrastructure Controller, on page 3](#)
- [About ACI/APIC Interfaces, on page 5](#)
- [Mixing the NX-OS Style CLI and the APIC GUI, on page 6](#)
- [Configuration Validation, on page 8](#)

Recommended Settings for the Cisco Application Policy Infrastructure Controller

We recommend the following settings for the Cisco Application Policy Infrastructure Controller (Cisco APIC):

Table 2: Recommended Settings for the Cisco APIC

Navigation Path	Property	Value	Description
System > System Settings > Fabric Wide Setting	Enforce Subnet Check	Put a check in the box.	This feature enforces subnet checks at the VRF instance level, when the Cisco Application Centric Infrastructure (Cisco ACI) learns the IP address as an endpoint from the data plane. Although the subnet check scope is the VRF instance, this feature can be enabled and disabled only globally under the fabric-wide setting policy. You cannot enable this option only in one VRF instance. If you put a check in the box for this option, the fabric will not learn IP addresses from a subnet other than the one configured on the bridge domain. This feature prevents the fabric from learning endpoint information in this scenario.
System > System Settings > Endpoint Controls	IP Aging Policy	Enabled	The IP aging policy tracks and ages unused IP addresses on an endpoint. Tracking is performed by using the endpoint retention policy, which is configured for the bridge domain to send ARP requests (for IPv4) and neighbor solicitations (for IPv6) at 75% of the local endpoint aging interval. When no response is received from an IP address, that IP address is aged out.

Navigation Path	Property	Value	Description
Fabric > External Access Policies > Policies > Global > MCP Instance Policy default	Admin State	Enabled	This enables the Mis-cabling Protocol (MCP)
	Controls: Enable MCP PDU per VLAN	Put a check in the box.	MCP detects other types of loops that can be caused by various issues, such as misconfiguration, that LLDP and STP cannot discover. This option enables MCP to send packets on a per-EPG basis.

About ACI/APIC Interfaces

The single point of management within the Cisco Application Centric Infrastructure (ACI) architecture is known as the Application Policy Infrastructure Controller (APIC). This controller provides access to all configuration, management, monitoring, and health functions. Having a centralized controller with an application programming interface (API) means that all functions configured or accessed through the fabric can be approached through the following interfaces:

- APIC GUI

The APIC GUI is a browser-based graphical interface to the APIC that communicates internally with the APIC engine by exchanging REST API messages. It includes two modes:

- Formerly called Advanced Mode, now simply the APIC GUI—Used for large scale configurations, deployments, and operations; enables granular policy controls such as in switch profiles, interface profiles, policy groups, or access entity profiles (AEPs) for automating mass fabric configuration and deployment.
- Formerly Basic Mode—Up to release 3.1(x), but now removed, this was a simple interface to enable common workflows, the GUI operational mode enables administrators to get started easily with ACI with a minimal knowledge of the object model. The simplified GUI allows the configuration of leaf ports and tenants without the need to configure advanced policies.

For more information about the APIC GUI, see *Cisco APIC Getting Started Guide, Release 3.x* and *Cisco APIC Basic Configuration Guide, Release 3.x*.

- NX-OS Style CLI—The NX-OS style Command-Line Interface (CLI) can be used for APIC configuration, deployment, and operation. It is organized in a hierarchy of command modes with EXEC mode as the root, containing a tree of configuration submodes beginning with global configuration mode. The commands available to you depend on the mode you are in.

For important guidelines to use both the NX-OS style CLI and the APIC GUI to configure Cisco APIC, see [Mixing the NX-OS Style CLI and the APIC GUI, on page 6](#).

For more information about the NX-OS style CLI, see *Cisco APIC NX-OS Style Command-Line Interface Configuration Guide*.

- APIC REST API—The REST API is responsible for accepting configuration, as well as providing access to management functions for the controller. This interface is a crucial component for the GUI and CLI, and also provides a touch point for automation tools, provisioning scripts and third party monitoring and management tools.

The APIC REST API is a programmatic interface that uses REST architecture. The API accepts and returns HTTP (not enabled by default) or HTTPS messages that contain JavaScript Object Notation (JSON) or Extensible Markup Language (XML) documents. You can use any programming language to generate the messages and the JSON or XML documents that contain the API methods or MO descriptions.

For more information about the REST API, see the *Cisco APIC REST API Configuration Guide*.

Mixing the NX-OS Style CLI and the APIC GUI

Basic mode is deprecated since Cisco APIC Release 3.0(1). There is only one GUI as of that release.



Caution

Configurations done through the NX-OS style CLI are rendered in the APIC GUI. They can be seen, but sometimes may not be editable in the GUI. Also changes made in the APIC GUI may be seen in the NX-OS style CLI, but may only partially work. See the following examples:

- Do not mix the GUI and the CLI, when doing per-interface configuration on APIC. Configurations performed in the GUI, may only partially work in the NX-OS CLI.

For example, if you configure a switch port in the GUI at **Tenants > *tenant-name* > Application Profiles > *application-profile-name* > Application EPGs > *EPG-name* > Static Ports > Deploy Static EPG on PC, VPC, or Interface**

Then you use the show running-config command in the NX-OS style CLI, you receive output such as:

```
leaf 102
interface ethernet 1/15
switchport trunk allowed vlan 201 tenant t1 application ap1 epg ep1
exit
exit
```

If you use these commands to configure a static port in the NX-OS style CLI, the following error occurs:

```
apic1(config)# leaf 102
apic1(config-leaf)# interface ethernet 1/15
apic1(config-leaf-if)# switchport trunk allowed vlan 201 tenant t1 application ap1 epg
ep1
No vlan-domain associated to node 102 interface ethernet1/15 encap vlan-201
```

This occurs because the CLI has validations that are not performed by the APIC GUI. For the commands from the show running-config command to function in the NX-OS CLI, a vlan-domain must have been previously configured. The order of configuration is not enforced in the GUI.

For the steps to remove such objects, see *Troubleshooting Unwanted _ui_ Objects* in the *APIC Troubleshooting Guide*.

About the Modes of Configuring Layer 3 External Connectivity

Because APIC supports multiple user interfaces (UIs) for configuration, the potential exists for unintended interactions when you create a configuration with one UI and later modify the configuration with another UI. This section describes considerations for configuring Layer 3 external connectivity with the APIC NX-OS style CLI, when you may also be using other APIC user interfaces.

When you configure Layer 3 external connectivity with the APIC NX-OS style CLI, you have the choice of two modes:

- Implicit mode, a simpler mode, is not compatible with the APIC GUI or the REST API.
- Named (or Explicit) mode is compatible with the APIC GUI and the REST API.

In either case, the configuration should be considered read-only in the incompatible UI.

How the Modes Differ

In both modes, the configuration settings are defined within an internal container object, the "L3 Outside" (or "L3Out"), which is an instance of the **L3extOut** class in the API. The main difference between the two modes is in the naming of this container object instance:

- Implicit mode—the naming of the container is implicit and does not appear in the CLI commands. The CLI creates and maintains these objects internally.
- Named mode—the naming is provided by the user. CLI commands in the Named Mode have an additional **L3Out** field. To configure the named L3Out correctly and avoid faults, the user is expected to understand the API object model for external Layer 3 configuration.



Note Except for the procedures in the *Configuring Layer 3 External Connectivity Using the Named Mode* section, this guide describes Implicit mode procedures.

Guidelines and Restrictions

- In the same APIC instance, both modes can be used together for configuring Layer 3 external connectivity with the following restriction: The Layer 3 external connectivity configuration for a given combination of tenant, VRF, and leaf can be done only through one mode.
- For a given tenant VRF, the policy domain where the External-L3 EPG can be placed can be in either the Named mode or in the Implicit mode. The recommended configuration method is to use only one mode for a given tenant VRF combination across all the nodes where the given tenant VRF is deployed for Layer 3 external connectivity. The modes can be different across different tenants or different VRFs and no restrictions apply.
- In some cases, an incoming configuration to a Cisco APIC cluster will be validated against inconsistencies, where the validations involve externally-visible configurations (northbound traffic through the L3Outs). An Invalid Configuration error message will appear for those situations where the configuration is invalid.
- The external Layer 3 features are supported in both configuration modes, with the following exception:
 - Route-peering and Route Health Injection (RHI) with a L4-L7 Service Appliance is supported only in the Named mode. The Named mode should be used across all border leaf switches for the tenant VRF where route-peering is involved.

- Layer 3 external network objects (l3extOut) created using the Implicit mode CLI procedures are identified by names starting with “_ui_” and are marked as read-only in the GUI. The CLI partitions these external-L3 networks by function, such as interfaces, protocols, route-map, and EPG. Configuration modifications performed through the REST API can break this structure, preventing further modification through the CLI.

For the steps to remove such objects, see *Troubleshooting Unwanted _ui_ Objects* in the *APIC Troubleshooting Guide*.

Configuration Validation

When the administrator enters a configuration in the Cisco Application Policy Infrastructure Controller (Cisco APIC), the Cisco APIC performs checks to make sure that the configuration is valid, which is known as validation. If the configuration is accepted, but it conflicts with other previous configurations, Cisco APIC or the leaf switches might raise faults. The amount of checks performed by the Cisco APIC before accepting a configuration varies depending on the release. Newer releases have been enhanced to perform more checks before the configuration is accepted instead of only raising faults asynchronously.

The release with the greatest amount of changes in terms of additional validations is the Cisco APIC release 2.3. Cisco APIC release 3.0 further enhances validations at the VRF instance level. As an example, in Cisco APIC release 2.3, for the same VRF instance and the same L3Out, you can define multiple Switch Virtual Interface (SVI) logical interface profiles for the same SVI (encap) with different IP addresses. You can define IP address 10.10.10.1/24 on path node1, port 1/41, VLAN (encap) 10, and IP address 10.10.10.2/24 for path node1, port 1/43, VLAN (encap) 10.

This results in only one IP address being used for SVI 10 on the leaf switch despite the fact that you configured multiple IP addresses, and depending on which IP address is used as the next hop for routing or whether you have IGP configured, the configuration might function properly.

Starting with Cisco APIC release 3.0, the above configuration would not be accepted, because even if in the Cisco Application Centric Infrastructure (Cisco ACI) object model the SVI is defined per path (logical interface profile), a given VRF instance on a given leaf switch can only have one IP address for an SVI and potentially a secondary IP address. Several other validations were also introduced in Cisco APIC release 3.0.

The objective of these validations is to reduce or eliminate configuration errors by informing the user of the errors at the configuration time instead of accepting the configuration and raising faults asynchronously.

As a result of these improvements, if you POST a configuration that was incorrect, but was considered valid prior to the 2.3 release, this POST would not result in the configuration being posted and the Cisco APIC will return an error message.

There might be existing Cisco APIC deployments that are functioning correctly with versions prior to Cisco APIC release 2.3 despite the fact that the configurations might not be valid. To reduce the impact of a firmware upgrade in such scenarios, after you upgrade to the 2.3 release or later, the Cisco APIC relaxes the validation checks on existing configurations.

Cisco APIC also offers the option to import an existing configuration with the "Best Effort" mode instead of the "Atomic" mode. This option offers the ability to accept a configuration even if there are portions that are not valid. The Cisco APIC pushes the valid portions of the configuration and ignores the portions that are not consistent with the validation. For the inconsistent portions, the Cisco APIC issues an error message that is visible when you use the following command:

```
show snapshot jobs import_job
```




CHAPTER 3

First Time Setup Wizard

This chapter contains the following sections:


- [About the First Time Setup Wizard, on page 9](#)

About the First Time Setup Wizard

Use the First Time Setup wizard to set up your Cisco APIC for the first time.

- You can access the First Time Setup wizard when it automatically appears the first time you log into your Cisco APIC through the GUI.
- For Cisco APIC Releases 4.2(3) and later, you can also access the First Time Setup wizard when you



click the System Tools icon () in the upper right corner of the Cisco APIC GUI window, then select **What's New in APIC_release_number**.

The **Welcome to APIC** window appears, providing information on the new features that are part of this particular release.

To access the First Time Setup wizard, click **Begin First Time Setup** or **Review First Time Setup** at the bottom right of the window. The **Let's Configure the Basics** window appears, with links to the individual pages that you can use to set up your Cisco APIC.

When you have completed the initial setup that includes at least one BGP route reflector, the **Proceed to Summary** button is enabled. Click this button to view summary tiles of the configuration. Additional tiles appear under the heading **You Might Want To...** These additional topics are optional but recommended.

The following sections provide more information for each of the first-time setup pages available from this window.

Fabric Membership

Use the **Fabric Membership** window to register the leaf and spine switches detected by the ACI fabric. You can also manually add leaf and spine switches to the fabric using the serial number listed on the box.




Note We recommend registering at least two leaf switches and two spine switches. You must register at least one leaf switch and one spine switch in order to proceed through the First Time Setup wizard.

The **Fabric Membership** window contains two sections:

- **Discovered:** This section provides information on newly-discovered but unregistered switches. These nodes will have a node ID of 0 and will have no IP address.
- **Registered:** This section provides information on all of the registered switches in your ACI fabric.

You can register a switch using either of these methods:

- If the switch is shown in the **Discovered** section, click the **Register** button next to that switch to open the **Create Fabric Node Member** window. Note that the **Pod ID** and **Serial Number** fields will be automatically populated in the **Create Fabric Node Member** window in this case.
- If the switch is not shown in the **Discovered** section, click the Action icon (), then select **Create Fabric Node Member** from the drop-down list.

In the **Create Fabric Node Member** window, enter the following information:

Field	Setting
Pod ID	Identify the pod where the node is located.
Serial Number	Required: Enter the serial number of the switch.
Node ID	<p>Required: Enter a number greater than 100. The first 100 IDs are reserved for APIC appliance nodes.</p> <p>Note We recommend that leaf nodes and spine nodes be numbered differently. For example, number leaves in the 100 range (such as 101, 102) and number spines in the 200 range (such as 201, 202).</p> <p>Note After the node ID is assigned, it cannot be updated. After the node has been added to the Registered Nodes tab table, you can update the node name by right-clicking the table row and choosing Edit Node and Rack Name.</p>
Switch Name	The node name, such as leaf1 or spine3.

Field	Setting
Node Type	<p>Choose the assigned node role. The options are:</p> <ul style="list-style-type: none"> • leaf <p>Check one of the following boxes if applicable:</p> <ul style="list-style-type: none"> • Is Remote • Is Virtual • Is Tier-2 Leaf <ul style="list-style-type: none"> • spine <p>Check the following box if applicable:</p> <ul style="list-style-type: none"> • Is Virtual <ul style="list-style-type: none"> • unknown

Click **Submit** when you have completed the information in the **Create Fabric Node Member** window, then click **Continue** in the **Fabric Membership** to continue to the next window in the First Time Setup wizard.

Management

Use the **Out Of Band Management** window to configure the management interface IP address for leaf switches, spine switches, and APIC nodes to connect to the Out of Band (OOB) network. Select several nodes to begin assigning IP addresses to them.



Note The First Time Setup wizard helps with configuring nodes that have not already been configured for Out of Band management.

Type the address of the gateway in the following fields and an IP address is automatically suggested for each discovered node:

- **IPv4 Gateway:** The IPv4 default gateway address for communication to external networks using out-of-band management.
- **IPv6 Gateway:** The IPv6 default gateway address for communication to external networks using out-of-band management.

In the Filter by attributes area, you can filter the discovered nodes by their attributes. Click **Edit** if you want to change the nodes that you had selected to configure for Out of Band management.

Click **Save** to continue to the next window in the First Time Setup wizard.

vPC Pairs

Use the **Setup - vPC Pairs** window to explicitly configure member nodes of the group by using a Fabric policy node endpoint.

In the Filter by attributes area, you can filter the vPC pairs by their attributes. Click **Edit** if you want to change the vPC pairs that you had selected to configure the vPC pair.

In the vPC Pairs central pane, expand the **Actions** drop-down list, and choose **Create vPC Leaf Switch Pair**, **Delete vPC Leaf Switch Pair**, **Download All**, or **Open in Object Store Browser**.

Click **Save** to continue to the next window in the First Time Setup wizard.

BGP

Use the **BGP** window to configure ACI fabric route reflectors, which use multiprotocol BGP (MP-BGP) to distribute external routes within the fabric. Once you have enabled the route reflectors in the ACI fabric, you can configure connectivity to external networks.



Note Select spine switches to configure as route reflectors. You must configure at least one route reflector in order to proceed through the First Time Setup wizard. If you do not see any spine switches in the table in this window, verify that the switch is registered with the correct type or has been discovered by APIC.

In the **BGP** window, check the box next to the spine switches that you want to use as route reflectors and enter the ASN for this spine switch in the **Autonomous System Number** field. Click **Save and Continue** to continue to the next window in the First Time Setup wizard.

DNS

Use the **DNS** window to configure DNS servers and search domains to allow leaf switches, spine switches and APIC nodes to query DNS names. The OOB connection will be used for DNS communication.



Note The First Time Setup wizard configures DNS servers and DNS domains under the **default** DNS Policy.

To configure the DNS servers, click + in the DNS Servers area, then enter the following information:

- **Address:** Enter the provider address.
- **Preferred:** Check the check box if you want to have this address as the preferred provider.
- **Status:** Provides the status of the configuration request.

Click **Update**, then repeat this process to configure additional DNS servers, if necessary.

To configure the search domains, click + in the Search Domains area, then enter the following information:

- **Name:** Enter the domain name (cisco.com).
- **Default:** Check the check box to make this domain the default domain. You can have only one domain name as the default.
- **Status:** Provides the status of the configuration request.

Click **Update**, then repeat this process to configure additional search domains, if necessary.

To delete an entry either from the DNS Servers table or from the Search Domains table, select the entry that you would like to delete, then click the trash can icon in that table.

Click **Save and Continue** to continue to the next window in the First Time Setup wizard.

NTP

Use the **NTP** window to configure a timezone and assign NTP servers to synchronize leaf switches, spine switches, and APIC nodes to a valid time source. The OOB connection will be used for NTP communication.



Note The First Time Setup wizard configures servers under the **default** NTP Policy.

In the Display Format area, click **local** to display the date and time in a local time zone format, or click **utc** to display the date and time in the UTC time zone format. The default is **local**.

If you selected **local** above, in the Time Zone area, click the drop-down arrow to choose the time zone for your domain. You can also type in the drop down menu area to filter the drop down options. The default is **Coordinated Universal Time**.

To configure the NTP servers, click + in the NTP Servers area, then enter the following information:

- **Host Name/IPAddress:** Enter the host name and IP address of the NTP server.
- **Preferred:** If you are creating multiple providers, check the **Preferred** check box for the most reliable NTP source.
- **Status:** Provides the status of the configuration request.

Click **Update**, then repeat this process to configure additional NTP servers, if necessary.

To delete an entry from the NTP Servers table, select the entry that you would like to delete, then click the trash can icon in that table.

Click **Save and Continue** to continue to the next window in the First Time Setup wizard.

Proxy

Use the **Proxy** window to configure the HTTP or HTTPS proxy policy. When configured, some of the Cisco Cloud Application Policy Infrastructure Controller (APIC) features, mainly those that need internet access such as Cisco Intersight connectivity, send the traffic through the HTTP or HTTPS proxy. For more information, see the *Cisco APIC System Management Configuration Guide*.

Global Configurations

Use the **Global Configurations** window to configure certain areas, which we recommend as best practices during the first time set up of your Cisco Application Centric Infrastructure (ACI) fabric. Click **Okay, Got it!** when you are ready to configure these areas:

- [Subnet Check, on page 14](#)
- [Domain Validation, on page 14](#)
- [Intermediate System to Intermediate System for redistributed routes, on page 14](#)

- [IP Aging Administrative State, on page 14](#)
- [Rogue Endpoint Control, on page 15](#)
- [COOP Group Policy, on page 15](#)



Note Some settings in this window are configurable after the First Time Setup, such as the Subnet Check and Domain Validation settings, which can be configured in the **Fabric Wide Setting Policy** page (**System > System Settings > Fabric-Wide Settings**). However, configuring those settings after the First Time Setup might cause issues with other existing configurations. For example, enabling the **Enforce Subnet Check** and **Enforce Domain Validation** settings in the **Fabric Wide Setting Policy** page could break a configured L3Out connection without the proper policy chain in place for the interface or for a statically-assigned port to an EPG.

Subnet Check

This feature disables IP address learning outside of subnets configured in a VRF instance, for all other VRF instances.

This feature enforces subnet checks at the VRF instance level, when the Cisco ACI learns the IP address as an endpoint from the data plane. If you put a check in the box for this option, the fabric will not learn IP addresses from a subnet other than the one configured on the bridge domain. This feature prevents the fabric from learning endpoint information in this scenario.

Put a check in the box next to **Enforce** to enable the subnet check feature, which we highly recommend.

Domain Validation

This feature enforces a validation check if a static path is added but no domain is associated to an EPG.

When enabled, a validation check is performed when a static path is added to an EPG, to determine if the path is part of a domain that is associated with the EPG. The scope of this policy is fabric-wide. After configuration, a policy is pushed to each leaf switch as it comes up.

Put a check in the box next to **Enforce** to enable the domain validation feature, which is highly recommended.

Intermediate System to Intermediate System for redistributed routes

This is the IS-IS metric that is used for all imported routes into IS-IS. Configuring a metric lower than 64 (max) with this option, such as 63, allows Cisco ACI switches to prefer routes from stable spines until the routing convergence is achieved on a new spine.

Enter the appropriate value in the **IS-IS metric** field.

IP Aging Administrative State

Enabling this policy allows Cisco ACI to track each IP address individually and age out unused addresses efficiently. Otherwise, unused IP addresses remain learned until the base MAC address ages out. This does not affect remote endpoints.

When enabled, the IP aging policy ages unused IP addresses on an endpoint. In this situation, the IP aging policy sends ARP requests (for IPv4) and neighbor solicitations (for IPv6) to track IP addresses on endpoints. If no response is given, the policy ages the unused IP addresses.

Following are the options for this field:

- **Disabled:** The default setting. Cisco APIC disregards the IP aging policy.
- **Enabled:** Cisco APIC observes the IP aging policy.

We highly recommend enabling this feature.

Rogue Endpoint Control

A rogue endpoint can attack leaf switches through frequently, repeatedly injecting packets on different leaf switch ports and changing 802.1Q tags (emulating endpoint moves), resulting in IP and MAC addresses being learned rapidly in different EPGs and ports. Misconfigurations can also cause frequent IP and MAC address changes (moves).

The rogue endpoint control feature addresses this vulnerability. Enabling this policy allows Cisco ACI to detect and delete unauthorized endpoints.

Following are the options for this field:

- **Disabled:** The default setting. Cisco APIC disregards the rogue endpoint control policy.
- **Enabled:** Cisco APIC observes the rogue endpoint control policy.

We highly recommend that you enable this feature.

Additional settings for rogue endpoint control, such as **Rogue EP Detection Interval**, **Rogue EP Detection Multiplication Factor**, and **Hold Interval**, are available through the **Endpoint Controls** panel. To access the **Endpoint Controls** panel, on the menu bar, click **System > System Settings > Endpoint Controls**, then click the **Rogue EP Control** tab.

Following are the valid and default settings for the fields in the **Rogue EP Control** tab in the **Endpoint Controls** window:

- **Rogue EP Detection Interval:** Valid values are from 0 to 65535 seconds. Default value is 60.
- **Rogue EP Detection Multiplication Factor:** Valid values are from 2 to 65535. Default value is 4.
- **Hold Interval:** In the 5.2(1) and 5.2(2) releases, the valid values are from 1800 to 3600 seconds. Beginning with the 5.2(3) release, the valid values are from 300 to 3600 seconds. The default value is 1800.

COOP Group Policy

Council of Oracle Protocol (COOP) is used to communicate the mapping information (location and identity) to the spine proxy. A leaf switch forwards endpoint address information to the spine switch 'Oracle' using Zero Message Queue (ZMQ). COOP running on the spine nodes will ensure all spine nodes maintain a consistent copy of endpoints and location information in the mapping database.

COOP protocol supports two ZMQ authentication modes:

- **Compatible Type:** The default setting. COOP accepts both MD5 authenticated and non-authenticated ZMQ connections for message transportation.



Note The Cisco APIC manages the token used as MD5 password for COOP. This token is automatically rotated by Cisco APIC every hour. This token cannot be displayed.

- **Strict Type:** COOP allows MD5 authenticated ZMQ connections only.

We highly recommend the **Strict Type** setting for the COOP group policy.



CHAPTER 4

User Access, Authentication, and Accounting

This chapter contains the following sections:

- [Access Rights Workflow Dependencies, on page 17](#)
- [User Access, Authorization, and Accounting, on page 17](#)
- [Login Domains, on page 20](#)
- [Creating Login Domain Using the GUI, on page 20](#)
- [Creating a Provider , on page 21](#)
- [Configuring a Local User, on page 25](#)
- [Configuring a Remote User, on page 27](#)
- [Configuring Windows Server 2012 LDAP for APIC Access with Cisco AVPair, on page 33](#)
- [Configuring APIC for LDAP Access, on page 35](#)
- [Changing the Default Behavior for Remote Users with Missing or Bad Cisco AV Pairs, on page 36](#)
- [About Signature-Based Transactions, on page 36](#)
- [Accounting, on page 43](#)
- [Routed Connectivity to External Networks as a Shared Service Billing and Statistics, on page 44](#)

Access Rights Workflow Dependencies

The Cisco Application Centric Infrastructure (ACI) RBAC rules enable or restrict access to some or all of the fabric. For example, in order to configure a leaf switch for bare metal server access, the logged in administrator must have rights to the `infra` domain. By default, a tenant administrator does not have rights to the `infra` domain. In this case, a tenant administrator who plans to use a bare metal server connected to a leaf switch could not complete all the necessary steps to do so. The tenant administrator would have to coordinate with a fabric administrator who has rights to the `infra` domain. The fabric administrator would set up the switch configuration policies that the tenant administrator would use to deploy an application policy that uses the bare metal server attached to an ACI leaf switch.

User Access, Authorization, and Accounting

Application Policy Infrastructure Controller (APIC) policies manage the authentication, authorization, and accounting (AAA) functions of the Cisco Application Centric Infrastructure (ACI) fabric. The combination of user privileges, roles, and domains with access rights inheritance enables administrators to configure AAA functions at the managed object level in a granular fashion. These configurations can be implemented using the REST API, the CLI, or the GUI.



Note There is a known limitation where you cannot have more than 32 characters for the login domain name. In addition, the combined number of characters for the login domain name and the user name cannot exceed 64 characters.

Multiple Tenant Support

A core Application Policy Infrastructure Controller (APIC) internal data access control system provides multitenant isolation and prevents information privacy from being compromised across tenants. Read/write restrictions prevent any tenant from seeing any other tenant's configuration, statistics, faults, or event data. Unless the administrator assigns permissions to do so, tenants are restricted from reading fabric configuration, policies, statistics, faults, or events.

User Access: Roles, Privileges, and Security Domains

The APIC provides access according to a user's role through role-based access control (RBAC). An Cisco Application Centric Infrastructure (ACI) fabric user is associated with the following:

- A predefined or custom role, which is a set of one or more privileges assigned to a user
- A set of privileges, which determine the managed objects (MOs) to which the user has access
- For each role, a privilege type: no access, read-only, or read-write
- One or more security domain tags that identify the portions of the management information tree (MIT) that a user can access

Roles and Privileges

A privilege controls access to a particular function within the system. The ACI fabric manages access privileges at the managed object (MO) level. Every object holds a list of the privileges that can read from it and a list of the privileges that can write to it. All objects that correspond to a particular function will have the privilege for that function in its read or write list. Because an object might correspond to additional functions, its lists might contain multiple privileges. When a user is assigned a role that contains a privilege, the user is given read access to the associated objects whose read list specifies read access, and write access to those whose write list specifies write access.

As an example, 'fabric-equipment' is a privilege that controls access to all objects that correspond to equipment in the physical fabric. An object corresponding to equipment in the physical fabric, such as 'eqptBoard,' will have 'fabric-equipment' in its list of privileges. The 'eqptBoard' object allows read-only access for the 'fabric-equipment' privilege. When a user is assigned a role such as 'fabric-admin' that contains the privilege 'fabric-equipment,' the user will have access to those equipment objects, including read-only access to the 'eqptBoard' object.



Note Some roles contain other roles. For example, '-admin' roles such as tenant-admin, fabric-admin, access-admin are groupings of roles with the same base name. For example, 'access-admin' is a grouping of 'access-connectivity', 'access-equipment', 'access-protocol', and 'access-qos.' Similarly, tenant-admin is a grouping of roles with a 'tenant' base, and fabric-admin is a grouping of roles with a 'fabric' base.

The 'admin' role contains all privileges.

For more details about roles and privileges see [APIC Roles and Privileges Matrix](#).

Security Domains

A security domain is a tag associated with a certain subtree in the ACI MIT object hierarchy. For example, the default tenant "common" has a domain tag `common`. Similarly, the special domain tag `all` includes the entire MIT object tree. An administrator can assign custom domain tags to the MIT object hierarchy. For example, an administrator could assign the "solar" domain tag to the tenant named solar. Within the MIT, only certain objects can be tagged as security domains. For example, a tenant can be tagged as a security domain but objects within a tenant cannot.



Note Security Domain password strength parameters can be configured by creating **Custom Conditions** or by selecting **Any Three Conditions** that are provided.

Creating a user and assigning a role to that user does not enable access rights. It is necessary to also assign the user to one or more security domains. By default, the ACI fabric includes two special pre-created domains:

- `All`—allows access to the entire MIT
- `Infra`— allows access to fabric infrastructure objects/subtrees, such as fabric access policies



Note For read operations to the managed objects that a user's credentials do not allow, a "DN/Class Not Found" error is returned, not "DN/Class Unauthorized to read." For write operations to a managed object that a user's credentials do not allow, an HTTP 401 Unauthorized error is returned. In the GUI, actions that a user's credentials do not allow, either they are not presented, or they are grayed out.

A set of predefined managed object classes can be associated with domains. These classes should not have overlapping containment. Examples of classes that support domain association are as follows:

- Layer 2 and Layer 3 network managed objects
- Network profiles (such as physical, Layer 2, Layer 3, management)
- QoS policies

When an object that can be associated with a domain is created, the user must assign domain(s) to the object within the limits of the user's access rights. Domain assignment can be modified at any time.

If a virtual machine management (VMM) domain is tagged as a security domain, the users contained in the security domain can access the correspondingly tagged VMM domain. For example, if a tenant named solar

is tagged with the security domain called sun and a VMM domain is also tagged with the security domain called sun, then users in the solar tenant can access the VMM domain according to their access rights.

Login Domains

A login domain defines the authentication domain for a user. Login domains can be set to the Local, LDAP, RADIUS, TACACS+, DUO, SAML, RSA, or OAuth 2 authentication mechanisms. When accessing the system from REST, the CLI, or the GUI, the APIC enables the user to select the correct authentication domain.

For example, in the REST scenario, the username is prefixed with a string so that the full login username looks as follows:

```
apic:<domain>\<username>
```

If accessing the system from the GUI, the APIC offers a drop-down list of domains for the user to select. If no `apic: domain` is specified, the default authentication domain servers are used to look up the username.

Starting in ACI version 1.0(2x), the login domain fallback of the APIC defaults local. If the default authentication is set to a non-local method and the console authentication method is also set to a non-local method and both non-local methods do not automatically fall back to local authentication, the APIC can still be accessed via local authentication.

To access the APIC fallback local authentication, use the following strings:

- From the GUI, use `apic:fallback\username`.
- From the REST API, use `apic#fallback\username`.



Note Do not change the fallback login domain. Doing so could result in being locked out of the system.

Creating Login Domain Using the GUI

Before you begin

- The ACI fabric is installed, Application Policy Infrastructure Controllers (APICs) are online, and the APIC cluster is formed and healthy.
- The login domain name, realm, and remote server provider group are available to define the authentication domain for the user.

Procedure

Step 1

In the APIC, create Login Domain.

- On the menu bar, choose **Admin > Authentication > AAA > Policy** tab.
- In the **Properties** pane, click the **Actions** icon > **Create Login Domain**. You can also create a login domain by clicking + displayed against **Login Domains**.

c) In the **Create Login Domain** pane, specify the following:

- The user configured domain name.
- Description of the login domain.
- The realm to verify the identity of an entity (person or device) accessing the fabric devices. The options available in the **Realm** drop-down list are discussed here:
 1. For Release 4.2(x) and earlier, choose a security method from Local, LDAP, RADIUS, TACACS+, RSA, or SAML for processing authentication requests.
 2. For Release 5.0(x) and later, choose a security method from DUO Proxy LDAP, DUO Proxy Radius, LDAP, RADIUS, TACACS+, RSA, SAML, OAuth 2, or Local for processing authentication requests.

Note

If LDAP, RADIUS, or TACACS+ is specified as the default security method and the associated provider group specified in this dialog is not available to provide authentication during a user login, fallback local authentication is not executed by the APIC server unless is specifically configured to do so.

- A RADIUS provider group for a group of remote servers supporting the RADIUS protocol for authentication.
- A TACACS+ provider group for a group of remote servers supporting the TACACS+ protocol for authentication.
- An LDAP provider group for a group of remote servers supporting the LDAP protocol for authentication.
- A RSA provider group for a group of remote servers supporting the RSA protocol for authentication.
- A SAML provider group for a group of remote servers supporting the SAML protocol for authentication.
- An OAuth 2 provider group for a group of remote servers supporting the OAuth 2 protocol for authentication.

Step 2 Click **Submit**.

Creating a Provider

Use this procedure to create a provider for the authentication/authorization protocols.

Before you begin

The relevant prerequisites before creating a provider for an authentication/authorization protocol is discussed under the relevant protocol sections.

Procedure

-
- Step 1** On the menu bar, choose **Admin > AAA**.
- Step 2** In the Navigation pane, choose **Authentication**.
- Step 3** In the Work pane, choose **Providers**.
- Step 4** Click **Actions > Create Provider**.
- Step 5** In the **Create Provider** screen that is displayed, enter the **Hostname/IP Address**, **Description**, and choose a **Realm** from the drop-down list. The choices available for **Realm** are:
- RADIUS
 - TACACS+
 - LDAP
 - SAML
 - RSA
 - OAuth 2

The options for configuring a provider are dynamic and change as per the selected **Realm**. The options available for each **Realm** are discussed in detail, in the subsequent steps.

- Step 6** (Optional) Applicable only for RADIUS: Choose a **Realm Subtype**. The options are **Default** or **Duo**. Specify the following:
- Password for the RADIUS server; enter the password again for confirmation.
 - Click **Select Reachability EPG** to choose an endpoint group.
 - The service port number for RADIUS. The range is 1 to 65535. The default value is 1812.
 - The authentication protocol, options are **PAP**, **CHAP**, **MS-CHAP**. You will see this option only if you chose **Default** for the **Realm Subtype**.
 - The timeout for communication with a RADIUS server. The range is from 0 to 60 seconds. The default is 5 seconds (Realm Subtype: Default); default is 30 seconds (Realm Subtype:Duo).
 - The number of retries when contacting the RADIUS endpoint.
 - Select the **Enabled** check box to enable periodic server monitoring; enter a user name and password for the same.

This step is for RADIUS provider configuration. You can now proceed to step 12.

- Step 7** (Optional step; applicable only for TACACS+) Specify the following:
- Password for the TACACS+ server; enter the password again for confirmation.
 - Click **Select Reachability EPG** to choose an endpoint group.
 - The service port number for TACACS+. The range is 1 to 65535. The default value is 49.
 - The authentication protocol, options are—PAP, CHAP, MS-CHAP.

- The timeout for communication with a TACACS+ server. The range is from 0 to 60 seconds. The default is 5 seconds.
- The number of retries when contacting the TACACS+ endpoint.
- Select the **Enabled** check box to enable periodic server monitoring; enter a user name and password for the same.

This step is for TACACS+ provider configuration. You can now proceed to step 12.

Step 8

(Optional step; applicable only for LDAP) Choose a Realm Subtype, options are—**Default** or **Duo**. Specify the following:

- The Root Distinguished Name (DN) of the LDAP directory.
- The LDAP Base DN, which is the container name and path in the LDAP server where the APIC searches for the remote user account. This is where the password is validated. Filter is used to locate the attribute that the APIC requests to use for the *Cisco AVPair*.
- Password for the LDAP server; enter the password again for confirmation.
- The service port number for LDAP. The range is 1 to 65535. The default value is 389.
- Click **Select Reachability EPG** to choose an endpoint group.
- The timeout for communication with an LDAP server. The range is from 0 to 60 seconds. The default is 30 seconds.
- The number of retries when contacting the LDAP endpoint.
- Select the **Enable** check box to enable SSL.
- SSL Certificate Validation Level. The options are:
 - Permissive—A debugging knob to help diagnose DUO LDAP SSL Certificate issues.
 - Strict—A level that should be used when in production.
- LDAP Attribute.
- Authentication Method. The options are:
 - LDAP Bind
 - Password Compare
- Filter Type. Filters are a key element in defining the criteria used to identify entries in search requests. Example: (cn=*), which means any entry that contains one or more *cn* values. The options are:
 - Default
 - Microsoft Active Directory
 - Custom
- LDAP Filter. This field is auto-filled based on the selected Filter Type (unless you have chosen the Custom option Filter Type). If you have chosen Default, the filter is `cn=Suserid`; if you have chosen Microsoft Active Directory, the filter is `sAMAccountName=Suserid`.

- Select the **Enabled** check box to enable periodic server monitoring; enter a user name and password for the same.

This step is for LDAP provider configuration. You can now proceed to step 12.

Step 9 (Optional step; applicable only for RSA) Specify the following:

- Password for the RSA server; enter the password again for confirmation.
- Click **Select Reachability EPG** to choose an endpoint group.
- The service port number for RSA. The range is 1 to 65535. The default value is 1812.
- The timeout for communication with a RSA server. The range is from 0 to 60 seconds. The default is 5 seconds.
- The number of retries when contacting the RSA endpoint.
- Select the **Enabled** check box to enable periodic server monitoring; enter a user name and password for the same.

This step is for RSA provider configuration. You can now proceed to step 12.

Step 10 (Optional step; applicable only for SAML) Specify the following:

- **Identity Provider (IdP)**. The options are—ADFS, OKTA, PING IDENTITY.
- **Metadata URL provided by IDP**.
In case of ADFS, **Metadata URL provided by IDP** is of the format, *https://<FQDN of ADFS>/FederationMetadata/2007-06/FederationMetadata.xml*.
For Ping ID, copy the metadata URL link from the configuration section of the Ping ID server (under the SAML application).
- Enter the **IdP Entity ID** for the SAML-based service.
- Click **Select Certificate Authority** to select a certificate authority if IdP is signed by a private CA.
- GUI Redirect Banner. This can be a URL or a message. This information is displayed before the user is redirected to the Identity Provider login page for authentication.
- The timeout for communication with a SAML server. The range is from 0 to 60 seconds. The default is 5 seconds.
- Select the Signature Algorithm from the drop down list.
- Put a check in the **Enabled** check box, to enable all/any of these— Encrypted SAML Assertions, Assertions in SAML Response Signed, SAML Auth Requests Signed, SAML Response Message Signed.

This step is for SAML provider configuration. You can now proceed to step 12.

Step 11 (Optional step; applicable only for OAuth 2) Specify the following:

- Client ID—Client identifier of the APIC application on IdP.
- Client Secret for the APIC application. Enter the client secret again for confirmation.
- Username Claim. Username attribute in the token. Example: email, sub.

- Scope. List of OAuth 2 scopes. Example: "openid profile". To receive user group information, add the corresponding scope configured in the IdP provider. Example: "openid profile groups".
- Choose to **Enable** or **Disable** the OIDC Protocol.
- Put a check in the **Enabled** check box to Verify Token Signature.
- JWKS Endpoint. The JSON Web Key Sets (JWKS) to verify the token. This field is displayed only if you have enabled Verify Token Signature.
- Authorization Endpoint. The IdP endpoint authorization URL. Get the authorization endpoint from the IdP server. This field is displayed only when the OIDC protocol is disabled.
- Token Endpoint. The IdP endpoint token URL. Get the token endpoint from the IdP server. This field is displayed only when the OIDC protocol is disabled.
- Issuer URL. Get the issuer URL from the IdP server. This field is displayed only when the OIDC protocol is enabled.
- Click **Select Certificate Authority** to select a certificate authority if IdP is signed by a private CA.
- Click **Select Reachability EPG** to choose an endpoint group.
- The timeout for communication with an OAuth 2 server. The range is from 0 to 60 seconds. The default is 5 seconds.
- GUI Redirect Banner. This can be a URL or a message. This information is displayed before the user is redirected to the Identity Provider login page for authentication.

This step is for OAuth 2 provider configuration. You can now proceed to step 12.

Step 12 Click **Save**.

Configuring a Local User

In the initial configuration script, the admin account is configured and the admin is the only user when the system starts. The APIC supports a granular, role-based access control system where user accounts can be created with various roles including non-admin users with fewer privileges.

Configuring a Local User Using the GUI

Before you begin

- The ACI fabric is installed, APIC controllers are online, and the APIC cluster is formed and healthy.
- As appropriate, the security domain(s) that the user will access are defined. For example, if the new user account will be restricted to accessing a tenant, the tenant domain is tagged accordingly.
- An APIC user account is available that will enable the following:
 - Creating the TACACS+ provider.

- Creating the local user account in the target security domain(s). If the target domain is `all`, the login account used to create the new local user must be a fabric-wide administrator that has access to `all`. If the target domain is a tenant, the login account used to create the new local user must be a tenant administrator that has full read write access rights to the target tenant domain.

Procedure

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

Step 2 In the **Navigation** pane, click **Users**.

In the **Work** pane, verify that you are in the **Local Users** tab.

Step 3 In the **Work** pane, click the task icon drop-down list and select **Create Local User**.

Step 4 In the **STEP 1 > User Identity** dialog box, perform the following actions:

a) In the **Login ID** field, add an ID.

The login ID must meet the following guidelines:

- Must be unique within APIC.
- Must begin with a letter.
- Can contain between 1 and 32 characters.
- Can include alphanumeric characters, underscores, dashes, and dots.

After creating a user account, you cannot change the login ID. You must delete the user account and create a new one.

b) In the **Password** field, enter the password.

At the time a user sets their password, the Cisco APIC validates it against the following criteria:

- Minimum password length is 8 characters.
- Maximum password length is 64 characters.
- Has fewer than three consecutive repeated characters.
- Must have characters from at least three of the following characters types: lowercase, uppercase, digit, symbol.
- Does not use easily guessed passwords.
- Cannot be the username or the reverse of the username.
- Cannot be any variation of `cisco`, `isco` or any permutation of these characters or variants obtained by changing the capitalization of letters therein.

c) In the **Confirm Password** field, confirm the password.

d) (Optional) For Certificate based authentication, in the **User Certificate Attribute** field, enter the user identity from the authentication certificate.

e) Click **Next**.

- Step 5** You can activate or deactivate the user account by using the **Account Status** control, and you can set an expiration date by using the **Account Expires** control.
- Step 6** In the **STEP 2 > Security** dialog box, under **Security Domain**, choose the desired security domain for the user, and click **Next**.
- Step 7** In the **STEP 3 > Roles** dialog box, perform the following actions:
- Click the + to associate the user with a domain.
 - From the drop-down lists, choose a **Role Name** and a **Role Privilege Type** for the user.
 - click **Update**
- You can provide read-only or read/write privileges.
- Step 8** click **Finish**
-

Configuring SSH Public Key Authentication Using the GUI

Before you begin

- Create a local user account in the target security domain(s). If the target domain is `all`, the login account used to create the new local user must be a fabric-wide administrator that has access to `all`. If the target domain is a tenant, the login account used to create the new local user must be a tenant administrator that has full read write access rights to the target tenant domain.
- Generate a public key using the Unix command **ssh-keygen**.
The default login domain must be set to **local**

Procedure

- Step 1** On the menu bar, choose **ADMIN > Users** and confirm you are in the **Local Users** tab.
- Step 2** In the **Navigation** pane, click the name of the user that you previously created.
- Step 3** In the **Work** pane, expand the **SSH Keys** table, and insert the following information:
- In the **Name** field, enter a name for the key.
 - In the **Key** field, insert the public key previously created. Click **Update**.

Note

To create the SSH Private Key File for downloading to a remote location then in the menu bar, expand **Firmware > Download Tasks**.

Configuring a Remote User

Instead of configuring local users, you can point the APIC at the centralized enterprise credential datacenter. The APIC supports Lightweight Directory Access Protocol (LDAP), active directory, RADIUS, and TACACS+.



Note When an APIC is in minority (disconnected from the cluster), remote logins can fail because the ACI is a distributed system and the user information is distributed across APICS. Local logins, however, continue to work because they are local to the APIC.

Starting with the 3.1(1) release, **Server Monitoring** can be configured through RADIUS, TACACS+, LDAP, and RSA to determine whether the respective AAA servers are alive or not. Server monitoring feature uses the respective protocol login to check for server aliveness. For example, a LDAP server will use ldap login and a Radius server will use radius login with server monitoring to determine server aliveness.

To configure a remote user authenticated through an external authentication provider, you must meet the following prerequisites:

- The DNS configuration should have already been resolved with the hostname of the RADIUS server.
- You must configure the management subnet.

AV Pair on the External Authentication Server

The Cisco APIC requires that an administrator configure a Cisco AV Pair on an external authentication server. The Cisco AV pair specifies the APIC required RBAC roles and privileges for the user. The Cisco AV Pair format is the same for RADIUS, LDAP, or TACACS+.

To configure a Cisco AV Pair on an external authentication server, an administrator adds a Cisco AV pair to the existing user record. The Cisco AV pair format is as follows:

```
shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,
domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2
shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,
domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2(16003)
```

Starting with Cisco APIC release 2.1, if no UNIX ID is provided in AV Pair, the APIC allocates the unique UNIX user ID internally.



Note The APIC Cisco AV-pair format is compatible and can co-exist with other Cisco AV-pair formats. APIC will pick up the first matching AV-pair from all the AV-pairs.

Starting with release 3.1(x), the AV Pair shell:domains=all/admin allows you to assign Read-only privileges to users and provide them access to the switches and run commands.

The APIC supports the following regexes:

```
shell:domains\s*[:]\s*((\\S+?/\\S*?/\\S*?) (, \\S+?/\\S*?/\\S*?) {0,31}) (\\ (\\d+\\)) $
shell:domains\s*[:]\s*((\\S+?/\\S*?/\\S*?) (, \\S+?/\\S*?/\\S*?) {0,31}) $
```

Examples:

- Example 1: A Cisco AV Pair that contains a single Security domain with only writeRoles:

```
shell:domains=domainA/writeRole1|writeRole2/
```

- Example 2: A Cisco AV Pair that contains a single Security domain with only readRoles:

```
shell:domains=domainA//readRole1|readRole2
```



Note The "/" character is a separator between writeRoles and readRoles per Security domain and is required even if only one type of role is to be used.

The Cisco AVpair string is case sensitive. Although a fault may not be seen, using mismatching cases for the domain name or roles could lead to unexpected privileges being given.

An example configuration for an open RADIUS server (/etc/raddb/users) is as follows:

```
aaa-network-admin Cleartext-Password := "<password>"
Cisco-avpair = "shell:domains = all/aaa/read-all(16001)"
```

Configuring an AV Pair on the External Authentication Server

The numerical value within the parentheses in the attribute/value (AV) pair string is used as the UNIX user ID of the user who is logged in using Secure Shell (SSH) or Telnet.



Note Beginning with the 5.3(1) release, telnet is not supported.

Procedure

Configure an AV pair on the external authentication server.

The Cisco AV pair definition is as follows (Cisco supports AV pairs with and without UNIX user IDs specified):

Example:

```
shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2
shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2(8101)
```

These are the boost regexes supported by APIC:

```
uid_regex("shell:domains\\s*[:=]\\s*(\\S+?/\\S*?/\\S*?)(,\\S+?/\\S*?/\\S*?){0,31})(\\d+\\S*)$");
regex("shell:domains\\s*[:=]\\s*(\\S+?/\\S*?/\\S*?)(,\\S+?/\\S*?/\\S*?){0,31}$");
```

The following is an example:

```
shell:domains = coke/tenant-admin/read-all,pepsi//read-all(16001)
```

Configuring APIC for TACACS+ Access

Before you begin

- The Cisco Application Centric Infrastructure (ACI) fabric is installed, Application Policy Infrastructure Controllers (APICs) are online, and the APIC cluster is formed and healthy.
- The TACACS+ server host name or IP address, port, and key are available.
- The APIC management endpoint group is available.

Procedure

Step 1 In the APIC, create the **TACACS+ Provider**.

- a) On the menu bar, choose **Admin > Authentication > TACACS** tab.
- b) Click the **Actions** icon > **Create TACACS+ Provider**.
- c) In the displayed pop-up window, specify the TACACS+ host name (or IP address), description, port, authorization protocol, key, timeout, retries, management endpoint group, and select if server monitoring needs to be enabled or not.

Note

If the APIC is configured for in-band management connectivity, out-of-band management does not work for authentication. With the APIC release 2.1(1x), you can set a global toggle between In-band and out-of-band as the default management connectivity between the APIC server and other external management devices.

For toggling in-band or out-of-band management in the APIC GUI:

- Prior to Release 2.2(1x): In the **Navigation** pane, choose **Fabric > Fabric Policies > Global Policies > Connectivity Preferences**. In the **Work Pane** select either **inband** or **ooband**.
- For Release 2.2(x) and 2.3(x): In the **Navigation** pane, choose **Fabric > Fabric Policies > Global Policies > APIC Connectivity Preferences**. In the **Work Pane** select either **inband** or **ooband**.
- For Release 3.0(1x) or later: In the **Navigation** pane, choose **System > System Settings > APIC Connectivity Preferences**. In the **Work Pane** select either **inband** or **ooband**.

Step 2 Create the **Login Domain** for TACACS+.

For the detailed procedure, see [Creating Login Domain Using the GUI, on page 20](#).

What to do next

This completes the APIC TACACS+ configuration steps. Next, if a RADIUS server will also be used, configure the APIC for RADIUS. If only a TACACS+ server will be used, go to the ACS server configuration topic below.

Configuring APIC for RADIUS Access

Before you begin

- The ACI fabric is installed, Application Policy Infrastructure Controllers (APICs) are online, and the APIC cluster is formed and healthy.
- The RADIUS server host name or IP address, port, authorization protocol, and key are available.
- The APIC management endpoint group is available.

Procedure

Step 1 In the APIC, create the RADIUS provider.

- a) On the menu bar, choose **Admin > Authentication > RADIUS** tab.
- b) Click the **Actions** icon > **Create RADIUS Provider**.
- c) In the displayed pop-up window, specify the RADIUS host name (or IP address), description, port, protocol, key, timeout, retries, management endpoint group, and select if server monitoring needs to be enabled or not. .

Note

If the APIC is configured for in-band management connectivity, out-of-band management does not work for authentication. With the APIC release 2.1(1x), you can set a global toggle between In-band and out-of-band as the default management connectivity between the APIC server and other external management devices.

For toggling in-band or out-of-band management in the APIC GUI:

- Prior to Release 2.2(1x): In the **Navigation** pane, choose **Fabric > Fabric Policies > Global Policies > Connectivity Preferences**. In the **Work Pane** select either **inband** or **ooband**.
- For Release 2.2(x) and 2.3(x): In the **Navigation** pane, choose **Fabric > Fabric Policies > Global Policies > APIC Connectivity Preferences**. In the **Work Pane** select either **inband** or **ooband**.
- For Release 3.0(1x) or later: In the **Navigation** pane, choose **System > System Settings > APIC Connectivity Preferences**. In the **Work Pane** select either **inband** or **ooband**.

Step 2 Create the **Login Domain** for TACACS+.

For the detailed procedure, see [Creating Login Domain Using the GUI, on page 20](#).

What to do next

This completes the APIC RADIUS configuration steps. Next, configure the RADIUS server.

Configuring a Cisco Secure Access Control Server for RADIUS and TACACS+ Access to the APIC

Before you begin

- The Cisco Secure Access Control Server (ACS) version 5.5 is installed and online.



Note ACS v5.5 was used to document these steps. Other versions of ACS might support this task but the GUI procedures might vary accordingly.

- The Cisco Application Policy Infrastructure Controller (Cisco APIC) RADIUS or TACACS+ keys are available (or keys for both if both will be configured).
- The APICs are installed and online; the APIC cluster is formed and healthy.
- The RADIUS or TACACS+ port, authorization protocol, and key are available.

Procedure

- Step 1** Log in to the ACS server to configure the APIC as a client.
- a) Navigate to **Network Resources > Network Devices Groups > Network Devices and AAA Clients**.
 - b) Specify the client name, the APIC in-band IP address, select the TACACS+ or RADIUS (or both) authentication options.
- Note**
If the only RADIUS or TACACS+ authentication is needed, select only the needed option.
- c) Specify the authentication details such as Shared Secret (key), and port as appropriate for the authentication option(s).
- Note**
The **Shared Secret(s)** must match the APIC **Provider** key(s).
- Step 2** Create the Identity Group.
- a) Navigate to **Users and Identity Stores > Internal Groups** option.
 - b) Specify the **Name**, and **Parent Group** as appropriate.
- Step 3** Map users to the Identity Group.
- a) In the **Navigation** pane, click the **Users and Identity Stores > Internal Identity Stores > Users** option.
 - b) Specify the user **Name**, and **Identity Group** as appropriate.
- Step 4** Create the Policy Element.
- a) Navigate to the **Policy Elements** option.
 - b) For RADIUS, specify the Authorization and Permissions > Network Access > Authorization Profiles **Name**. For TACACS+, specify the Authorization and Permissions > Device Administration > Shell Profile **Name** as appropriate.

- c) For RADIUS, specify the **Attribute** as `cisco-av-pair`, **Type** as string, and the **Value** as `shell:domains = <domain>/<role>/,<domain>// role` as appropriate. For TACACS+, specify the **Attribute** as `cisco-av-pair`, **Requirement** as Mandatory, and the **Value** as `shell:domains = <domain>/<role>/,<domain>// role` as appropriate.

The syntax of the **Value** field determines whether write privileges are granted:

- For read/write privileges, the syntax is `shell:domains = <domain>/<role>/.`
- For read-only privileges, the syntax is `shell:domains = <domain>// <role>.`

For example, if the `cisco-av-pair` has a value of `shell:domains = solar/admin/,common// read-all`, then `solar` is the security domain, `admin` is the role that gives write privileges to this user in the security domain called `solar`, `common` is the tenant common, and `read-all` is the role with read privileges that gives this user read privileges to all of the tenant common.

Step 5 Create a service selection rule.

- a) For RADIUS, create a service selection rule to associate the Identity Group with the Policy Element by navigating to **Access Policies > Default Device Network Access Identity > Authorization** and specifying the rule **Name**, **Status**, and **Conditions** as appropriate, and **Add** the `Internal Users:UserIdentityGroup` in `ALL Groups:<identity group name>`.
- b) For TACACS+, create a service selection rule to associate the Identity Group with the Shell Profile by navigating to **Access Policies > Default Device Admin Identity > Authorization**. Specify the rule **Name**, **Conditions**, and **Select** the **Shell Profile** as appropriate.

What to do next

Use the newly created RADIUS and TACACS+ users to log in to the APIC. Verify that the users have access to the correct APIC security domain according to the assigned RBAC roles and privileges. The users should not have access to items that have not been explicitly permitted. Read and write access rights should match those configured for that user.

Configuring Windows Server 2012 LDAP for APIC Access with Cisco AVPair

Before you begin

- First, configure the LDAP server, then configure the Cisco Application Policy Infrastructure Controller (Cisco APIC) for LDAP access.
- The Microsoft Windows Server 2012 is installed and online.
- The Microsoft Windows Server 2012 Server Manager ADSI Edit tool is installed. To install ADSI Edit, follow the instructions in the Windows Server 2012 Server Manager help.
- `CiscoAVPair` attribute specifications: Common Name = **CiscoAVPair**, LDAP Display Name = **CiscoAVPair**, Unique X500 Object ID = 1.3.6.1.4.1.9.22.1, Description = **CiscoAVPair**, Syntax = **Case Sensitive String**.



Note For LDAP configurations, best practice is to use **CiscoAVPair** as the attribute string. If customer faces the issue using Object ID 1.3.6.1.4.1.9.22.1, an additional Object ID 1.3.6.1.4.1.9.2742.1-5 can also be used in the LDAP server.

- A Microsoft Windows Server 2012 user account is available that will enable the following:
 - Running ADSI Edit to add the `CiscoAVPair` attribute to the Active Directory (AD) Schema.
 - Configuring an Active Directory LDAP user to have `CiscoAVPair` attribute permissions.
- Port 636 is required for configuring LDAP integration with SSL/TLS.

Procedure

Step 1 Log in to an Active Directory (AD) server as a domain administrator.

Step 2 Add the `CiscoAVPair` attribute to the AD schema.

- a) Navigate to **Start > Run**, type `mmc` and press **Enter**.
The Microsoft Management Console (MMC) opens.
- b) Navigate to **File > Add/Remove Snap-in > Add**.
- c) In the **Add Standalone Snap-in** dialog box, select the **Active Directory Schema** and click **Add**.
The MMC **Console** opens.
- d) Right-click the **Attributes** folder, select the **Create Attribute** option.
The **Create New Attribute** dialog box opens.
- e) Enter `CiscoAVPair` for the **Common Name**, `CiscoAVPair` for the **LDAP Display Name**, `1.3.6.1.4.1.9.22.1` for the **Unique X500 Object ID**, and select **Case Sensitive String** for the **Syntax**.
- f) Click **OK** to save the attribute.

Step 3 Update the **User Properties** class to include the `CiscoAVPair` attribute.

- a) In the MMC **Console**, expand the **Classes** folder, right-click the **user** class, and choose **Properties**.
The **user Properties** dialog box opens.
- b) Click the **Attributes** tab, and click **Add** to open the **Select Schema Object** window.
- c) In the **Select a schema object:** list, choose `CiscoAVPair`, and click **Apply**.
- d) In the MMC **Console**, right-click the **Active Directory Schema**, and select **Reload the Schema**.

Step 4 Configure the `CiscoAVPair` attribute permissions.

Now that the LDAP includes the `CiscoAVPair` attributes, LDAP users need to be granted Cisco APIC permission by assigning them Cisco APIC RBAC roles.

- a) In the ADSI Edit dialog box, locate a user who needs access to the Cisco APIC.
- b) Right-click on the user name, and choose **Properties**.
The `<user> Properties` dialog box opens.
- c) Click the **Attribute Editor** tab, select the `CiscoAVPair` attribute, and enter the *Value* as `shell:domains = <domain>/<role>/, <domain>// role`.

For example, if the *CiscoAVPair* has a value of `shell:domains = solar/admin/,common//read-all(16001)`, then `solar` is the security domain, `admin` is the role for this user that gives write privileges to this user in the security domain called `solar`, `common` is the Cisco Application Centric Infrastructure (Cisco ACI) tenant `common`, and `read-all(16001)` is the role with read privileges that gives this user read privileges to all of the Cisco ACI tenant `common`.

- d) Click **OK** to save the changes and close the **<user> Properties** dialog box.

The LDAP server is configured to access the Cisco APIC.

What to do next

Configure the Cisco APIC for LDAP access.

Configuring APIC for LDAP Access

Before you begin

- The Cisco Application Centric Infrastructure (ACI) fabric is installed, Application Policy Infrastructure Controllers (APICs) are online, and the APIC cluster is formed and healthy.
- The LDAP server host name or IP address, port, bind DN, Base DN, and password are available.
- The APIC management endpoint group is available.

Procedure

Step 1 In the APIC, configure the LDAP Provider.

- a) On the menu bar, choose **Admin > Authentication > LDAP > Providers** tab.
- b) Click the **Actions** icon > **Create LDAP Provider**.
- c) In the displayed pop-up window, specify the LDAP host name (or IP address), port, bind DN, base DN, password, attribute, retries, timeout, SSL certificate validation level, filter type, management endpoint group, and select if server monitoring needs to be enabled or not. .

Note

- The bind DN is the string that the APIC uses to log in to the LDAP server. The APIC uses this account to validate the remote user attempting to log in. The base DN is the container name and path in the LDAP server where the APIC searches for the remote user account. This is where the password is validated. Filter is used to locate the attribute that the APIC requests to use for the *cisco-av-pair*. This contains the user authorization and assigned RBAC roles for use on the APIC. The APIC requests the attribute from the LDAP server.
- **Attribute** field—Enter one of the following:
 - For LDAP server configurations with a Cisco AVPair, enter **CiscoAVPair**.
 - For LDAP server configurations with an LDAP group map, enter **memberOf**.

- If the APIC is configured for in-band management connectivity, choosing an out-of-band management endpoint group for LDAP access does not take effect. Alternatively, an out-of-band over an in-band management endpoint group can connect a LDAP server, but requires configuring a static route for the LDAP server. The sample configuration procedures in this document use an APIC in-band management endpoint group.

Step 2 Create the **Login Domain** for LDAP.

For the detailed procedure, see [Creating Login Domain Using the GUI, on page 20](#).

What to do next

This completes the APIC LDAP configuration steps. Next, test the APIC LDAP login access.

Changing the Default Behavior for Remote Users with Missing or Bad Cisco AV Pairs

Procedure

Step 1 On the menu bar, choose **Admin > Authentication > AAA > Policy** tab.

Step 2 From the **Remote user login policy** drop-down list, choose **Assign Default Role**.

The default value is **No Login**. The **Assign Default Role** option assigns the minimal read-only privileges to users that have missing or bad Cisco AV Pairs. Bad AV Pairs are those AV Pairs that fail the parsing rules.

About Signature-Based Transactions

The APIC controllers in a Cisco ACI fabric offer different methods to authenticate users.

The primary authentication method uses a username and password and the APIC REST API returns an authentication token that can be used for future access to the APIC. This may be considered insecure in a situation where HTTPS is not available or enabled.

Another form of authentication that is offered utilizes a signature that is calculated for every transaction. The calculation of that signature uses a private key that must be kept secret in a secure location. When the APIC receives a request with a signature rather than a token, the APIC utilizes an X.509 certificate to verify the signature. In signature-based authentication, every transaction to the APIC must have a newly calculated signature. This is not a task that a user should do manually for each transaction. Ideally this function should be utilized by a script or an application that communicates with the APIC. This method is the most secure as it requires an attacker to crack the RSA/DSA key to forge or impersonate the user credentials.



Note Additionally, you must use HTTPS to prevent replay attacks.

Before you can use X.509 certificate-based signatures for authentication, verify that the following pre-requisite tasks are completed:

1. Create an X.509 certificate and private key using OpenSSL or a similar tool.
2. Create a local user on the APIC. (If a local user is already available, this task is optional).
3. Add the X.509 certificate to the local user on the APIC.

Guidelines and Limitations

Follow these guidelines and limitations:

- Local users are supported. Remote AAA users are not supported.
- The APIC GUI does not support the certificate authentication method.
- WebSockets and eventchannels do not work for X.509 requests.
- Certificates signed by a third party are not supported. Use a self-signed certificate.

Generating an X.509 Certificate and a Private Key

Procedure

Step 1 Enter an OpenSSL command to generate an X.509 certificate and private key.

Example:

```
$ openssl req -new -newkey rsa:1024 -days 36500 -nodes -x509 -keyout userabc.key -out userabc.crt -subj '/CN=User ABC/O=Cisco Systems/C=US'
```

Note

- Once the X.509 certificate is generated, it will be added to the users profile on the APIC, and it is used to verify signatures. The private key is used by the client to generate the signatures.
- The certificate contains a public key but not the private key. The public key is the primary information used by the APIC to verify the calculated signature. The private key is never stored on the APIC. You must keep it secret.

Step 2 Display the fields in the certificate using OpenSSL.

Example:

```
$ openssl x509 -text -in userabc.crt
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number:
```

```

c4:27:6c:4d:69:7c:d2:b6
Signature Algorithm: sha1WithRSAEncryption
Issuer: CN=User ABC, O=Cisco Systems, C=US
Validity
  Not Before: Jan 12 16:36:14 2015 GMT
  Not After : Dec 19 16:36:14 2114 GMT
Subject: CN=User ABC, O=Cisco Systems, C=US
Subject Public Key Info:
  Public Key Algorithm: rsaEncryption
  RSA Public Key: (1024 bit)
    Modulus (1024 bit):
      00:92:35:12:cd:2b:78:ef:9d:ca:0e:11:77:77:3a:
      99:d3:25:42:94:b5:3e:8a:32:55:ce:e9:21:2a:ff:
      e0:e4:22:58:6d:40:98:b1:0d:42:21:db:cd:44:26:
      50:77:e5:fa:b6:10:57:d1:ec:95:e9:86:d7:3c:99:
      ce:c4:7f:61:1d:3c:9e:ae:d8:88:be:80:a0:4a:90:
      d2:22:e9:1b:25:27:cd:7d:f3:a5:8f:cf:16:a8:e1:
      3a:3f:68:0b:9c:7c:cb:70:b9:c7:3f:e8:db:85:d8:
      98:f6:e3:70:4e:47:e2:59:03:49:01:83:8e:50:4a:
      5f:bc:35:d2:b1:07:be:ec:e1
    Exponent: 65537 (0x10001)
X509v3 extensions:
  X509v3 Subject Key Identifier:
    0B:E4:11:C7:23:46:10:4F:D1:10:4C:C1:58:C2:1E:18:E8:6D:85:34
  X509v3 Authority Key Identifier:
    keyid:0B:E4:11:C7:23:46:10:4F:D1:10:4C:C1:58:C2:1E:18:E8:6D:85:34
    DirName:/CN=User ABC/O=Cisco Systems/C=US
    serial:C4:27:6C:4D:69:7C:D2:B6

  X509v3 Basic Constraints:
    CA:TRUE
Signature Algorithm: sha1WithRSAEncryption
8f:c4:9f:84:06:30:59:0c:d2:8a:09:96:a2:69:3d:cf:ef:79:
91:ea:cd:ae:80:16:df:16:31:3b:69:89:f7:5a:24:1f:fd:9f:
d1:d9:b2:02:41:01:b9:e9:8d:da:a8:4c:1e:e5:9b:3e:1d:65:
84:ff:e8:ad:55:3e:90:a0:a2:fb:3e:3e:ef:c2:11:3d:1b:e6:
f4:5e:d2:92:e8:24:61:43:59:ec:ea:d2:bb:c9:9a:7a:04:91:
8e:91:bb:9d:33:d4:28:b5:13:ce:dc:fe:c3:e5:33:97:5d:37:
cc:5f:ad:af:5a:aa:f4:a3:a8:50:66:7d:f4:fb:78:72:9d:56:
91:2c
[snip]

```

Configuring a Local User

Creating a Local User and Adding a User Certificate Using the GUI

Procedure

- Step 1** On the menu bar, choose **ADMIN > AAA**.
- Step 2** In the **Navigation** pane, click **Users** and **Local Users** in the **Work** pane.
- Step 3** In the **Work** pane, verify that you in the **Local Users** tab.
The admin user is present by default
- Step 4** In the **Work** pane, click on task icon drop-down list and select **Create Local User**.

- Step 5** In the **Security** dialog box, choose the desired security domain for the user, and click **Next**.
- Step 6** In the **Roles** dialog box, click the radio buttons to choose the roles for your user, and click **Next**.
You can provide read-only or read/write privileges.
- Step 7** In the **User Identity** dialog box, perform the following actions:
- In the **Login ID** field, add an ID.
 - In the **Password** field, enter the password.
 - In the **Confirm Password** field, confirm the password.
 - (Optional) For Certificate based authentication, in the **User Certificate Attribute** field, enter the user identity from the authentication certificate.
 - Click **Finish**.
- Step 8** In the **Navigation** pane, click the name of the user that you created. In the **Work** pane, expand the + sign next to your user in the **Security Domains** area.
The access privileges for your user are displayed.
- Step 9** In the **Work** pane, in the **User Certificates** area, click the user certificates + sign, and in the **Create X509 Certificate** dialog box, perform the following actions:
- In the **Name** field, enter a certificate name.
 - In the **Data** field, enter the user certificate details.
 - Click **Submit**.
- The X509 certificate is created for the local user.

Creating a Local User Using Python SDK

Procedure

Create a local user.

Example:

```
#!/usr/bin/env python
from cobra.model.pol import Uni as PolUni
from cobra.model.aaa import UserEp as AaaUserEp
from cobra.model.aaa import User as AaaUser
from cobra.model.aaa import UserCert as AaaUserCert
from cobra.model.aaa import UserDomain as AaaUserDomain
from cobra.model.aaa import UserRole as AaaUserRole
from cobra.mit.access import MoDirectory
from cobra.mit.session import LoginSession
from cobra.internal.codec.jsoncodec import toJSONStr

APIC = 'http://10.10.10.1'
username = 'admin'
password = 'p@$w0rd'

session = LoginSession(APIC, username, password)
modir = MoDirectory(session)
modir.login()

def readFile(fileName=None, mode="r"):
```

```

if fileName is None:
    return ""
fileData = ""
with open(fileName, mode) as aFile:
    fileData = aFile.read()
return fileData

# Use a dictionary to define the domain and a list of tuples to define
# our aaaUserRoles (roleName, privType)
# This can further be abstracted by doing a query to get the valid
# roles, that is what the GUI does

userRoles = {'all': [
    ('aaa', 'writePriv'),
    ('access-admin', 'writePriv'),
    ('admin', 'writePriv'),
    ('fabric-admin', 'writePriv'),
    ('nw-svc-admin', 'writePriv'),
    ('ops', 'writePriv'),
    ('read-all', 'writePriv'),
    ('tenant-admin', 'writePriv'),
    ('tenant-ext-admin', 'writePriv'),
    ('vmm-admin', 'writePriv'),
    ],
}

uni = PolUni('') # '' is the Dn string for topRoot
aaaUserEp = AaaUserEp(uni)
aaaUser = AaaUser(aaaUserEp, 'userabc', firstName='Adam',
    email='userabc@cisco.com')

aaaUser.lastName = 'BC'
aaaUser.phone = '555-111-2222'
aaaUserCert = AaaUserCert(aaaUser, 'userabc.crt')
aaaUserCert.data = readFile("/tmp/userabc.crt")
# Now add each aaaUserRole to the aaaUserDomains which are added to the
# aaaUserCert
for domain,roles in userRoles.items():
    aaaUserDomain = AaaUserDomain(aaaUser, domain)
    for roleName, privType in roles:
        aaaUserRole = AaaUserRole(aaaUserDomain, roleName,
            privType=privType)
print toJSONStr(aaaUser, prettyPrint=True)

cr = ConfigRequest()
cr.addMo(aaaUser)
modir.commit(cr)
# End of Script to create a user

```

Using a Private Key to Calculate a Signature

Before you begin

You must have the following information available:

- HTTP method - GET, POST, DELETE
- REST API URI being requested, including any query options

- For POST requests, the actual payload being sent to the APIC
- The private key used to generate the X.509 certificate for the user
- The distinguished name for the user X.509 certificate on the APIC

Procedure

-
- Step 1** Concatenate the HTTP method, REST API URI, and payload together in this order and save them to a file. This concatenated data must be saved to a file for OpenSSL to calculate the signature. In this example, we use a filename of `payload.txt`. Remember that the private key is in a file called `userabc.key`.
- Example:**
- GET example:
- ```
GET http://10.10.10.1/api/class/fvTenant.json?rsp-subtree=children
```
- POST example:
- ```
POST http://10.10.10.1/api/mo/tn-test.json{"fvTenant": {"attributes": {"status": "deleted", "name": "test"}}
```
- Step 2** Verify that the `payload.txt` file contains the correct information. For example, using the GET example shown in the previous step:
- ```
GET http://10.10.10.1/api/class/fvTenant.json?rsp-subtree=children
```
- Your `payload.txt` file should contain only the following information:
- ```
GET/api/class/fvTenant.json?rsp-subtree=children
```
- Step 3** Verify that you didn't inadvertently create a new line when you created the payload file.
- Example:**
- ```
cat -e payload.txt
```
- Determine if there is a `$` symbol at the end of the output, similar to the following:
- ```
GET/api/class/fvTenant.json?rsp= subtree=children$
```
- If so, then that means that a new line was created when you created the payload file. To prevent creating a new line when generating the payload file, use a command similar to the following:
- ```
echo -n "GET/api/class/fvTenant.json?rsp-subtree=children" >payload.txt
```
- Step 4** Calculate a signature using the private key and the payload file using OpenSSL.
- Example:**
- ```
openssl dgst -sha256 -sign userabc.key payload.txt > payload_sig.bin
```
- The resulting file has the signature printed on multiple lines.
- Step 5** Convert the signature to base64 format:
- Example:**
- ```
openssl base64 -A -in payload_sig.bin -out payload_sig.base64
```
- Step 6** Strip the signature of the new lines using Bash.

**Example:**

```
$ tr -d '\n' < payload_sig.base64
P+OTqK0CeAZj17+Gute2R1Ww8OGgtzE0wsLlx8fIXX14V79Z17
Ou8IdJH9CB4W6CEvdICXqkv3KaQszCIC0+Bn07o3qF//BsIplZmYChD6gCX3f7q
IcjGX+R6HAqGeK7k97cNhXlWEoobFPe/oajtPjOu3tdOjhF/9ujG6Jv6Ro=
```

**Note**

This is the signature that will be sent to the APIC for this specific request. Other requests will require to have their own signatures calculated.

**Step 7** Place the signature inside a string to enable the APIC to verify the signature against the payload.

This complete signature is sent to the APIC as a cookie in the header of the request.

**Example:**

```
APIC-Request-Signature=P+OTqK0CeAZj17+Gute2R1Ww8OGgtzE0wsLlx8f
IXX14V79Z17Ou8IdJH9CB4W6CEvdICXqkv3KaQszCIC0+Bn07o3qF//BsIplZmYChD6gCX3f
7qIcjGX+R6HAqGeK7k97cNhXlWEoobFPe/oajtPjOu3tdOjhF/9ujG6Jv6Ro=;
APIC-Certificate-Algorithm=v1.0; APIC-Certificate-Fingerprint=fingerprint;
APIC-Certificate-DN=uni/userext/user-userabc/usercert-userabc.crt
```

**Note**

The DN used here must match the DN of the user certified object containing the x509 certificate in the next step.

**Step 8** Use the CertSession class in the Python SDK to communicate with an APIC using signatures.

The following script is an example of how to use the CertSession class in the ACI Python SDK to make requests to an APIC using signatures.

**Example:**

```
#!/usr/bin/env python
It is assumed the user has the X.509 certificate already added to
their local user configuration on the APIC
from cobra.mit.session import CertSession
from cobra.mit.access import MoDirectory

def readFile(fileName=None, mode="r"):
 if fileName is None:
 return ""
 fileData = ""
 with open(fileName, mode) as aFile:
 fileData = aFile.read()
 return fileData

pkey = readFile("/tmp/userabc.key")
csession = CertSession("https://ApicIPorHostname/",
 "uni/userext/user-userabc/usercert-userabc", pkey)

modir = MoDirectory(csession)
resp = modir.lookupByDn('uni/fabric')
print resp.dn
End of script
```

**Note**

The DN used in the earlier step must match the DN of the user certified object containing the x509 certificate in this step.

## Accounting

Cisco Application Centric Infrastructure (ACI) fabric accounting is handled by these two managed objects that are processed by the same mechanism as faults and events:

- The `aaaSessionLR` managed object tracks user account login and logout sessions on the Cisco Application Policy Infrastructure Controller (APIC) and switches, and token refresh. The Cisco ACI fabric session alert feature stores information such as the following:
  - Username
  - IP address initiating the session
  - Type (telnet, HTTPS, REST, and so on)



---

**Note** Beginning with the 5.3(1) release, telnet is not supported.

---

- Session time and length
- Token refresh: A user account login event generates a valid active token which is required in order for the user account to exercise its rights in the Cisco ACI fabric.



---

**Note** Token expiration is independent of login; a user could log out but the token expires according to the duration of the timer value it contains.

---

- The `aaaModLR` managed object tracks the changes users make to objects and when the changes occurred.
- If the AAA server is not pingable, it is marked unavailable and a fault is seen.

Both the `aaaSessionLR` and `aaaModLR` event logs are stored in Cisco APIC shards. After the data exceeds the pre-set storage allocation size, it overwrites records on a first-in first-out basis.



---

**Note** In the event of a destructive event such as a disk crash or a fire that destroys a Cisco APIC cluster node, the event logs are lost; event logs are not replicated across the cluster.

---

The `aaaModLR` and `aaaSessionLR` managed objects can be queried by class or by distinguished name (DN). A class query provides all the log records for the whole fabric. All `aaaModLR` records for the whole fabric are available from the GUI at the **Fabric > Inventory > POD > History > Audit Log** section. The Cisco APIC GUI **History > Audit Log** options enable viewing event logs for a specific object identified in the GUI.

The standard syslog, callhome, REST query, and CLI export mechanisms are fully supported for `aaaModLR` and `aaaSessionLR` managed object query data. There is no default policy to export this data.

There are no pre-configured queries in the Cisco APIC that report on aggregations of data across a set of objects or for the entire system. A fabric administrator can configure export policies that periodically export `aaaModLR` and `aaaSessionLR` query data to a syslog server. Exported data can be archived periodically and used to generate custom reports from portions of the system or across the entire set of system logs.

## Routed Connectivity to External Networks as a Shared Service Billing and Statistics

The Cisco Application Policy Infrastructure Controller (APIC) can be configured to collect byte count and packet count billing statistics from a port configured for routed connectivity to external networks as a shared service. The external networks are represented as external L3Out endpoint group (`l3extInstP` managed object) in Cisco Application Centric Infrastructure (ACI). Any EPG in any tenant can share an external L3Out EPG for routed connectivity to external networks. Billing statistics can be collected for each EPG in any tenant that uses an external L3Out EPG as a shared service. The leaf switch where the external L3Out EPG is provisioned forwards the billing statistics to the Cisco APIC where they are aggregated. Accounting policies can be configured to export these billing statistics periodically to a server.



## CHAPTER 5

# Management

---


This chapter contains the following sections:

- [Management Workflows, on page 45](#)
- [Adding Management Access, on page 47](#)
- [Exporting Tech Support, Statistics, and Core Files, on page 54](#)
- [Overview, on page 56](#)
- [Backing up, Restoring, and Rolling Back Controller Configuration, on page 62](#)
- [Using Syslog, on page 72](#)
- [Using Atomic Counters, on page 75](#)
- [Using SNMP, on page 78](#)
- [Using SPAN, on page 84](#)
- [Using Traceroute, on page 112](#)
- [acidiag Command, on page 114](#)

## Management Workflows

### ACI Management Access Workflows

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



## 1. Prerequisites

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

## 2. Configure the ACI Leaf Switch Access Ports

Choose which of these management access scenarios you will use:

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

### Suggested topics

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

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

## Adding Management Access

Configuring the external management instance profile under the management tenant for in-band has no effect on the protocols that are configured under the fabric-wide communication policies. The subnets and contracts specified under the external management instance profile do not affect HTTP/HTTPS or SSH/Telnet. Beginning with the 5.3(1) release, Telnet is not supported.

## Adding Management Access in the GUI

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

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

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

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

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

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



---

**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:
  - proto icmp, rate-limited, not configurable
  - tcp dpt: 22, rate-limited, not configurable
  - tcp dpt: 80, by default no listening process
  - tcp dpt: 443, default UI/API
  - tcp dpt: 4200, SSH access through web, by default no listening process

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

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

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

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

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

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

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

.

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

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

## Procedure

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

## Configuring In-Band Management Access Using the Cisco APIC GUI



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

---

## Procedure

- 
- Step 1** On the menu bar, choose **Fabric > Access Policies**.
- Step 2** In the **Navigation** pane, right-click **Interfaces** and choose **Configure Interface, PC and VPC**.
- Step 3** In the **Configure Interface, PC, and VPC** dialog box, to configure switch ports connected to Cisco Application Policy Infrastructure Controllers (APICs), perform the following actions:
- Click the large + icon next to the switch diagram to create a new profile and configure VLANs for the Cisco APIC.
  - From the **Switches** field drop-down list, check the check boxes for the switches to which the Cisco APICs are connected. (leaf1 and leaf2).
  - In the **Switch Profile Name** field, enter a name for the profile (apicConnectedLeaves).
  - Click the + icon to configure the ports.
  - Verify that in the **Interface Type** area, the **Individual** radio button is selected.
  - In the **Interfaces** field, enter the ports to which the Cisco APICs are connected.
  - In the **Interface Selector Name** field, enter the name of the port profile (apicConnectedPorts).

- 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:
- In the **Create Node Management Addresses** dialog box, in the **Policy Name** field, enter the policy name (apicInb).
  - In the **Nodes** field, **Select** column, check the check boxes for the nodes that will be part of this fabric (apic1, apic2, apic3).
  - In the **Config** field, check the **In-Band Addresses** check box.
  - In the **Node Range** fields, enter the range.
  - 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.
  - In the **In-Band IP Addresses** and **Gateway** fields, enter the IPv4 or IPv6 addresses as desired.
  - 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:
- In the **Create Node Management Addresses** dialog box, in the **Policy Name** field, enter the policy name (switchInb).
  - 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).
  - In the **Config** field, click the **In-Band Addresses** checkbox.
  - In the **Node Range** fields, enter the range.
  - 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.
  - In the **In-Band IP Addresses** and **Gateway** fields, enter the IPv4 or IPv6 addresses as desired.
  - Click **Submit**. In the **Confirm** dialog box, click **Yes**. The IP addresses for the leaf and spine switches are now configured.
- Step 12** In the **Navigation** pane, under **Node Management Addresses**, click the Cisco APIC policy name (apicInb) to verify the configurations. In the **Work** pane, the IP addresses assigned to various nodes are displayed.
- Step 13** In the **Navigation** pane, under **Node Management Addresses**, click the switches policy name (switchInb). In the **Work** pane, the IP addresses that are assigned to switches and the gateway addresses they are using are displayed.

**Note**

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

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



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

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

### Before you begin

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

## Procedure

- 
- Step 1** On the menu bar, choose **Tenants > mgmt**. In the **Navigation** pane, expand **Tenant mgmt**.
- Step 2** Right-click **Node Management Addresses**, and click **Create Node Management Addresses**.
- Step 3** In the **Create Node Management Addresses** dialog box, perform the following actions:
- In the **Policy Name** field, enter a policy name (switchOob).
  - In the **Nodes** field, check the check boxes next to the appropriate leaf and spine switches (leaf1, leaf2, spine1).
  - In the **Config** field, check the check box for **Out of-Band Addresses**.
- Note**  
The **Out-of-Band IP addresses** area is displayed.
- In the **Out-of-Band Management EPG** field, choose the EPG from the drop-down list (default).
  - In the **Out-Of-Band Gateway** field, enter the IP address and network mask for the external out-of-band management network.
  - 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:
- In the **Name** field, enter a name for the contract (oob-default).
  - Expand **Subjects**. In the **Create Contract Subject** dialog box, in the **Name** field, enter a subject name (oob-default).
  - 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**.
  - In the **Create Out-of-Band Contract** dialog box, click **Submit**.
- An out-of-band contract that can be applied to the out-of-band EPG is created.
- Step 8** In the **Navigation** pane, expand **Node Management EPGs > Out-of-Band EPG - default**.
- Step 9** In the **Work** pane, expand **Provided Out-of-Band Contracts**.
- Step 10** In the **OOB Contract** column, from the drop-down list, choose the out-of-band contract that you created (oob-default). Click **Update**, and click **Submit**.  
The contract is associated with the node management EPG.
- Step 11** In the **Navigation** pane, right-click **External EPG**, and click **Create External Management Entity Instance**.

- Step 12** In the **Create External Management Entity Instance** dialog box, perform the following actions:
- In the **Name** field, enter a name (oob-mgmt-ext).
  - 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.
  - In the **Subnets** field, enter the subnet address. Click **Submit**.  
Only the subnet addresses you choose here will be used to manage the switches. The subnet addresses that are not included cannot be used to manage the switches.

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

**Note**

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

---

## Exporting Tech Support, Statistics, and Core Files

### About Exporting Files

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

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

### File Export Guidelines and Restrictions

- HTTP export and the streaming API format is supported only with statistics information. Core and tech support data are not supported.
- The destination IP address for exported files cannot be an IPv6 address.
- Do not trigger tech support from more than five nodes simultaneously, especially if they are to be exported into the Cisco Application Policy Infrastructure Controller (APIC) or to an external server with insufficient bandwidth and compute resources.
- To collect tech support from all of the nodes in the fabric periodically, you must create multiple policies. Each policy must cover a subset of the nodes and should be scheduled to trigger in a staggered way (at least 30 minutes apart).

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

## Creating a Remote Location for Exporting Files

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

### Procedure

- 
- Step 1** In the menu bar, click **Admin**.
  - Step 2** In the submenu bar, click **Import/Export**.
  - Step 3** In the **Navigation** pane, expand **Export Policies**.
  - Step 4** Right-click **Remote Locations** and choose **Create Remote Path of a File**.
  - Step 5** In the **Create Remote Path of a File** dialog box, perform the following actions:
    - 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 Cisco Application Policy Infrastructure Controller (APIC) GUI. Use the following procedure to trigger a backup of your data.



---

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

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

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

---

### Procedure

- 
- Step 1** On the menu bar, choose **Admin > Import/Export**.
- Step 2** In the **Navigation** pane, right-click **Export Policies** and choose **Create Configuration Export Policy**. The **Create Configuration Export Policy** dialog appears.

- Step 3** Enter the appropriate values in the **Create Configuration Export Policy** dialog fields.  
For an explanation of a field, click the help (?) icon to display the help file.
- Step 4** After you finish entering values in the **Create Configuration Export Policy** dialog fields, click **Submit**.  
You have now created a backup. You can view this under the **Configuration** tab. The backup file will appear in the **Configuration** pane on the right.

**Note**

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

- Step 5** In the **Navigation** pane, choose **Export Policies > Configuration > policy\_name**.

- Step 6** In the **Work** pane, choose the **Operational > Job Status** tabs.

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

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

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

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

---

## Configuring an Import Policy Using the GUI

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

### Procedure

- 
- Step 1** On the menu bar, choose **ADMIN > Import/Export**.
- Step 2** In the navigation pane, right-click **Import Policies** and click **Create Configuration Import Policy**.  
The **Create Configuration Import Policy** dialog appears.
- Step 3** Enter the appropriate values in the **Create Configuration Import Policy** dialog fields.

**Note**

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

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

**Note**

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

## Encrypting Configuration Files Using the GUI

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

This section explains how to enable AES-256 encryption.

### Procedure

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

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

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

**Step 4** Click **SUBMIT**.

**Note**

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

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

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

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

**Note**

- Be sure that the **Fail Import if secure fields cannot be decrypted** check box is checked (which is the default selection) in your import and export policies. We highly recommend that you do not uncheck this box when you import configurations. If you uncheck this box, the system attempts to import all the fields. However, any fields that it cannot encrypt are blank/missing. As a result, you could lock yourself out of the system because the admin passwords could go blank/missing (if you lock yourself out of the system, refer to *Cisco APIC Troubleshooting Guide*). Unchecking the box launches a warning message. If the box is checked, there are security checks that prevent lockouts and the configuration does not import.

- When the **Enable Encryption** check box is unchecked (off), encryption is disabled and all exported configurations (exports) are missing the secure fields (such as passwords and certificates). When this box is checked (on), encryption is enabled and all exports show the secure fields.
- After enabling encryption, you cannot configure a passphrase when creating a new import or export policy. The passphrase you previously set is now global across all configurations in this box and across all tenants. If you export a configuration from this tab (you have configured a passphrase and enabled encryption) you get a complete backup file. If encryption is not enabled, you get a backup file with the secure properties removed. These backup files are useful when exporting to TAC support engineers, for example, because all the secure fields are missing. This is true for any secure properties in the configuration. There is also a clear option that clears the encryption key.

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

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

# Backing up, Restoring, and Rolling Back Controller Configuration

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

## Backing Up, Restoring, and Rolling Back Configuration Files Workflow

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

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



---

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

---

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

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

## About the fileRemotePath Object

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

- Hostname or IP

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




---

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

---

### Sample Configuration

The following is a sample configuration:

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

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

## Configuration Export to Controller

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




---

**Note** See the Snapshots section for more details.

---

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

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




---

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

---

## Scheduling Exports

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




---

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

---

## Troubleshooting

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

## Sample Configuration Using the NX-OS Style CLI

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

```

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

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

```

## Sample Configuration Using the GUI

The following is a sample configuration using the GUI:

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



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

### Sample Configuration Using REST API

The following is a sample configuration using the REST API:

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



---

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

---

## Configuration Import to Controller

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

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

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

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




---

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

---

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

- **importType**

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

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

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

### Troubleshooting

The following scenarios may need troubleshooting:

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

### Sample Configuration Using the NX-OS Style CLI

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

```

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

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

```

### Sample Configuration Using the GUI

The following is a sample configuration using the GUI:

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

### Sample Configuration Using the REST API

The following shows a sample configuration using the REST API:

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

## Snapshots

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

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

configSnapshot objects provide the following:

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

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

## Snapshot Manager Policy

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

You can also create a recurring snapshot.




---

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

---

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

### Troubleshooting

For troubleshooting, refer to the Connectivity Issues section.

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

```

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

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

```

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

```

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

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

```

### Snapshot Upload and Download Using the GUI

To upload a snapshot file to a remote location:

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

To download a snapshot file from a remote location:

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

### Snapshot Upload and Download Using the REST API

```

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

```

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

## Rollback

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

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



### Note

- The rollback feature only operates on snapshots.
- Remote archives are not supported directly. However, you can turn a remotely saved export into a snapshot using the snapshot manager policy (configSnapshotMgrP). For more information, see the [Snapshot Manager Policy, on page 68](#)
- The configRollbackP policy does not require a remote path configuration. If a remote path is provided, it will be ignored.

### Rollback Workflow

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

The MOs are handled as follows:

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

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

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

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

### Diff Tool

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

### Sample Configuration Using the NX-OS Style CLI

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

```
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](http://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/syslog/guide/aci_syslog/ACI_SysMsg.html).

---

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

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

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

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

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

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




---

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

---

## Creating a Syslog Destination and Destination Group

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



## Procedure

- Step 1** In the menu bar, click **Admin**.
- Step 2** In the submenu bar, click **External Data Collectors**.
- Step 3** In the **Navigation** pane, expand **Monitoring Destinations**.
- Step 4** Right-click **Syslog** and choose **Create Syslog Monitoring Destination Group**.
- Step 5** In the **Create Syslog Monitoring Destination Group** dialog box, perform the following actions:
- In the group and profile **Name** field, enter a name for the monitoring destination group and profile.
  - In the group and profile **Format** field, choose the format for Syslog messages.  
  
The default is **aci**, or the RFC 3164 compliant message format, but you can choose to set it to the NX-OS style format instead.
  - In the group and profile **Admin State** drop-down list, choose **enabled**.
  - To enable sending of syslog messages to a local file, choose **enabled** from the Local File Destination **Admin State** drop-down list and choose a minimum severity from the Local File Destination **Severity** drop-down list.  
  
The local file for receiving syslog messages is `/var/log/external/messages`.
  - To enable sending of syslog messages to the console, choose **enabled** from the Console Destination **Admin State** drop-down list and choose a minimum severity from the Console Destination **Severity** drop-down list.
  - Click **Next**.
  - In the **Create Remote Destinations** area, click **+** to add a remote destination.

### Caution

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

- Step 6** In the **Create Syslog Remote Destination** dialog box, perform the following actions:
- In the **Host** field, enter an IP address or a fully qualified domain name for the destination host.
  - (Optional) In the **Name** field, enter a name for the destination host.
  - In the **Admin State** field, click the **enabled** radio button.
  - (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.
  - 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 leaves.

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

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

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

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




---

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

---

Tenant atomic counters can provide the following:

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




---

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

---

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

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

## Atomic Counters Guidelines and Restrictions

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

endpoint-to-endpoint and EPG-to-EPG policies are supported. External policies are virtual routing and forwarding (VRF)-based, requiring learned IP addresses, and are supported.

- When the atomic counter source or destination is an endpoint, the endpoint must be dynamic and not static. Unlike a dynamic endpoint (fv:CEp), a static endpoint (fv:StCEp) does not have a child object (fv:RsCEpToPathEp) that is required by the atomic counter.
- In a transit topology, where leaf switches are not in full mesh with all spine switches, then leaf-to-leaf (TEP to TEP) counters do not work as expected.
- For leaf-to-leaf (TEP to TEP) atomic counters, once the number of tunnels increases the hardware limit, the system changes the mode from trail mode to path mode and the user is no longer presented with per-spine traffic.
- The atomic counter does not count spine proxy traffic.
- Packets dropped before entering the fabric or before being forwarded to a leaf port are ignored by atomic counters.
- Packets that are switched in the hypervisor (same Port Group and Host) are not counted.
- Atomic counters require an active fabric Network Time Protocol (NTP) policy.
- Atomic counters work for IPv6 sources and destinations, but configuring source and destination IP addresses across IPv4 and IPv6 addresses is not allowed.
- An atomic counter policy configured with fvCEp as the source or destination counts only the traffic that is from/to the MAC and IP addresses that are present in the fvCEp managed objects. If the fvCEp managed object has an empty IP address field, then all traffic to/from that MAC address would be counted regardless of the IP address. If the Cisco APIC has learned multiple IP addresses for an fvCEp, then traffic from only the one IP address in the fvCEp managed object itself is counted as previously stated. To configure an atomic counter policy to or from a specific IP address, use the fvIp managed object as the source or destination.
- If there is an fvIp behind an fvCEp, you must add fvIP-based policies and not fvCEp-based policies.
- Endpoint-to-endpoint atomic counter statistics are not reported for Layer 2 bridged traffic with IPv6 headers when the endpoints belong to the same EPG.
- For atomic counters to work for traffic flowing from an EPG or ESG to an L3Out EPG, configure the L3Out EPG with 0/1 and 128/1 to match all prefixes instead of 0/0.
- If your Cisco APIC has the traffic map mode set to "trail" 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 trail 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:
- In the **Name** field, enter a name for the policy.
  - 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).
  - choose or enter the identifying information for the traffic destination.
  - (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.
  - 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 3: 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 82 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 80](#) 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:
- In the **Name** field, enter an SNMP policy name.
  - In the **Admin State** field, select **Enabled**.
  - (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.
  - 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.
  - 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:
- In the **Client Group Policies** table, click the + icon to open the **Create SNMP Client Group Profile** dialog box.
  - In the **Name** field, enter an SNMP client group profile name.
  - From the **Associated Management EPG** drop-down list, choose the management EPG.
  - In the **Client Entries** table, click the + icon.
  - Enter a client's name in the **Name** field, enter the client's IP address in the **Address** field, and click **Update**.

**Note**

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

- Step 8** Click **OK**.
- Step 9** Click **Submit**.
- Step 10** Under **Pod Policies**, expand **Policy Groups** and choose a policy group or right-click **Policy Groups** and choose **Create POD Policy Group**.

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

- Step 11** In the pod policy group dialog box, perform the following actions:
- In the **Name** field, enter a pod policy group name.
  - From the **SNMP Policy** drop-down list, choose the SNMP policy that you configured and click **Submit**.
- Step 12** Under **Pod Policies**, expand **Profiles** and click **default**.
- Step 13** In the **Work pane**, from the **Fabric Policy Group** drop-down list, choose the pod policy group that you created.
- Step 14** Click **Submit**.
- Step 15** Click **OK**.

## Configuring an SNMP Trap Destination Using the GUI

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



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

### Procedure

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

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

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

- 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*.

## Multinode SPAN

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

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

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

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

## SPAN Guidelines and Restrictions



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

- Generation 1 switches are identified by the lack of a suffix, such as "EX", "FX", or "FX2," at the end of the switch name (for example, N9K-9312TX).
  - Generation 2 switches are identified with a suffix, such as "EX", "FX", or "FX2," at the end of the switch name.
- 
- The type of SPAN supported varies:
    - For generation 1 switches, tenant and access SPAN use the encapsulated remote extension of SPAN (ERSPAN) type I (Version 1 option in the Cisco Application Policy Infrastructure Controller (APIC) GUI).
    - For generation 2 switches, tenant and access SPAN use the encapsulated remote extension of SPAN (ERSPAN) type II (Version 2 option in the Cisco APIC GUI).
    - Fabric SPAN uses ERSPAN type II.
  - Beginning with the 5.2(3) release, ERSPAN supports IPv6 destinations.
  - The 6.0(3) release supports the Cisco N9K-C9808 switch, which has the following SPAN limitations:
    - Egress (transit (Tx)) SPAN is not supported.
    - SPAN on drop is not supported.

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

There are no limitations for Rx SPAN.

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




---

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

---

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

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

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

- EPG
- Routed outside (L3Out)
- When deploying the access SPAN source with an L3Out filter, ensure that the L3Out is also deployed on the matching interface:
  - If an L3Out is deployed on a port, a SPAN source must be deployed on the same port.

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

When an EPG filter is enabled within a SPAN session, ARP packets, which are sent out of the interface in the transit, or tx, direction, will not be spanned.

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

- When an access or fabric port-drop session and a global-drop session are configured, the access or fabric port-drop session takes the priority over the global-drop session.
- The number of filter entries supported in TCAM =  $(M * S1 * 1 + N * S2 * 2) + (S3 * 2)$ . This is applicable to rx SPAN or tx SPAN, separately. Currently, the maximum filter entries supported in tx or rx SPAN is 480 in each direction when following this formula (and assuming there are no other sources that are configured without filter-group association [means  $S3 = 0$ ] and with 16 port-ranges included). When the number of filter entries exceed the maximum number allowed, a fault will be raised. Note that you can specify Layer 4 port ranges in the filter entry. However, sixteen Layer 4 ports are programmed into the hardware as a single filter entry.

**Note**

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

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

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

When you configure SPAN for a given packet, SPAN is supported for the packet only once. If traffic is selected by SPAN in Rx for the first SSN, the traffic will not be selected by SPAN again in Tx for a second SSN. Thus, when the SPAN session ingress and egress 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.

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

## Configuring SPAN Using the GUI

### Configuring a Tenant SPAN Session Using the Cisco APIC GUI

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

#### Procedure

- 
- Step 1** In the menu bar, click **Tenants**.
  - Step 2** In the submenu bar, click the tenant that contains the source endpoint.
  - Step 3** In the **Navigation** pane, expand the tenant, expand **Policies > Troubleshooting > SPAN**.  
Two nodes appear under **SPAN**: **SPAN Destination Groups** and **SPAN Source Groups**.
  - Step 4** From the **Navigation** pane, right-click **SPAN Source Groups** and choose **Create SPAN Source Group**. The **Create SPAN Source Group** dialog appears.
  - Step 5** Enter the appropriate values in the required fields of the **Create SPAN Source Group** dialog box.
  - Step 6** Expand the **Create Sources** table to open the **Create SPAN Source** dialog box.
  - Step 7** Enter the appropriate values in the **Create SPAN Source** dialog box fields.
  - Step 8** When finished creating the SPAN source, click **OK**.

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

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

---

#### What to do next

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

## Configuring a SPAN Filter Group Using the APIC GUI

### Procedure

---

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

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

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

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

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

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

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

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

## Configuring an Access SPAN Policy Using the Cisco APIC GUI

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

### Procedure

- Step 1** In the menu bar, click on **Fabric > Access Policies**.
- Step 2** In the **Navigation** pane, expand **Policies > Troubleshooting > SPAN**.  
Three nodes appear under **SPAN**: **SPAN Source Groups**, **SPAN Filter Groups**, and **SPAN Destination Groups**.
- Step 3** Right-click **SPAN Source Groups** and choose **Create SPAN Source Group**.  
The **Create SPAN Source Group** dialog appears.
- Step 4** Enter the appropriate values in the **Create SPAN Source Group** dialog box fields.
- Step 5** Expand the **Create Sources** table to open the **Create SPAN Source** dialog box and enter the appropriate values in the required fields.
- Step 6** In the **Create SPAN Source** dialog box, expand **Add Source Access Paths** to specify the source path.  
The **Associate Source to Path** dialog box appears.
- Step 7** Enter the appropriate values in the **Associate Source to Path** dialog box fields.
- Step 8** When finished associating the source to a path, click **OK**.

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

**Step 9** When finished configuring the SPAN source, click **OK**.

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

**Step 10** When finished configuring the SPAN source group, click **Submit**.

---

#### What to do next

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

## Configuring a Fabric SPAN Policy Using the Cisco APIC GUI

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

### Procedure

---

**Step 1** In the menu bar, click on **Fabric > Fabric Policies**.

**Step 2** In the **Navigation** pane, expand **Policies > Troubleshooting > SPAN**.

Three nodes appear under **SPAN**: **SPAN Source Groups**, **SPAN Filter Groups**, and **SPAN Destination Groups**.

**Step 3** Right-click **SPAN Source Groups** and choose **Create SPAN Source Group**.  
The **Create SPAN Source Group** dialog appears.

**Step 4** Enter the appropriate values in the **Create SPAN Source Group** dialog box fields.

**Step 5** Expand the **Create Sources** table to open the **Create SPAN Source** dialog box.

**Step 6** Enter the appropriate values in the **Create SPAN Source** dialog box fields.

**Step 7** When finished, click **OK**.

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

**Step 8** When finished entering values in the **Create SPAN Source Group** dialog box fields, click **Submit**.

---

#### What to do next

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

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

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

## Procedure

- 
- Step 1** In the menu bar, click on **Fabric > Access Policies**.
- Step 2** In the **Navigation** pane, expand **Policies > Troubleshooting > SPAN**.  
Three nodes appear under **SPAN**: **SPAN Source Groups**, **SPAN Filter Groups**, and **SPAN Destination Groups**.
- Step 3** Right-click **SPAN Source Groups** and choose **Create SPAN Source Group**.  
The **Create SPAN Source Group** dialog appears.
- Step 4** Enter the appropriate values in the **Create SPAN Source Group** dialog box fields.
- Step 5** In the **Filter Group** field, select or create a filter group.  
See [Configuring a SPAN Filter Group Using the APIC GUI, on page 90](#) for more information.
- Step 6** Expand the **Create Sources** table to open the **Create SPAN Source** dialog box and perform the following actions:
- Enter a **Name** for the source policy.
  - Choose a **Direction** option for the traffic flow.
  - (Optional) Click to place a check mark in the **Span Drop Packets** check box. When checked, the SPAN-on-drop feature is enabled.
  - For external access, click **Routed Outside** in the **Type** field.
- Note**  
If **Routed Outside** is chosen for external access, then the **Name**, **Address**, and **Encap** fields appear to configure the **L3 Outside**.
- Expand **Add Source Access Paths** to specify the source path.  
The **Associate Source to Path** dialog box appears.
  - Enter the appropriate values in the **Associate Source to Path** dialog box fields.
  - When finished associating the source to a path, click **OK**.  
You return to the **Create SPAN Source** dialog box.
  - When finished configuring the SPAN source, click **OK**.  
You return to the **Create SPAN Source Group** dialog box.
- Step 7** When finished configuring the SPAN source group, click **Submit**.
- 

### What to do next

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

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

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

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

### Procedure

---

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

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

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

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

### Procedure

---

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

The destination group is created.

---

#### What to do next

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

## Configuring SPAN Using the NX-OS-Style CLI

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

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

#### Procedure

---

**Step 1**      **configure terminal**

Enters global configuration mode.

**Example:**

```
apic1# configure terminal
```

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

Creates an access monitoring session configuration.

**Example:**

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

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

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

**Example:**

```
apic1(config-monitor-access)# description "This is my SPAN session"
```

**Step 4**      **[no] destination interface ethernet *slot/port leaf node-id***

Specifies the destination interface. The destination interface cannot be a FEX port.

**Example:**

```
apic1(config-monitor-access)# destination interface ethernet 1/2 leaf 101
```

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

Specifies the source interface port or port range.

**Example:**

```
apic1(config-monitor-access)# source interface ethernet 1/2 leaf 101
```

**Step 6**      **drop enable**

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

**Example:**

```
apicl(config-monitor-access-source)# drop enable
```

**Step 7** [no] **direction** {rx | tx | both}

Specifies direction of traffic to be monitored. The direction can be configured independently for each source port range.

**Example:**

```
apicl(config-monitor-access-source)# direction tx
```

**Step 8** [no] **filter tenant** *tenant-name* **application** *application-name* **epg** *epg-name*

Filters traffic to be monitored. The filter can be configured independently for each source port range.

**Example:**

```
apicl(config-monitor-access-source)# filter tenant t1 application appl epg epg1
```

**Step 9** **exit**

Returns to access monitor session configuration mode.

**Example:**

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

**Step 10** [no] **destination interface port-channel** *port-channel-name-list* **leaf** *node-id*

Specifies the destination interface. The destination interface cannot be a FEX port.

**Note**

Beginning with Release 4.1(1), support is now available for having a static port-channel as the destination interface, as shown in the example command.

**Example:**

```
apicl(config-monitor-access)# destination interface port-channel pc1 leaf 101
```

**Step 11** [no] **source interface port-channel** *port-channel-name-list* **leaf** *node-id* [**fex** *fex-id*]

Specifies the source interface port channel.

(Enters the traffic direction and filter configuration, not shown here.)

**Example:**

```
apicl(config-monitor-access)# source interface port-channel pc5 leaf 101
```

**Step 12** [no] **filter tenant** *tenant-name* **l3out** *L3Out-name* **vlan** *interface-VLAN*

Filters traffic to be monitored. The filter can be configured independently for each source port range.

**Note**

Beginning with Release 4.1(1), you no longer have to specify the IP prefix when configuring L3Out interface filtering, as shown in the example.

**Example:**

```
apicl(config-monitor-access-source)# filter tenant t1 l3out l3out1 vlan 2820
```



**Step 13** [no] shutdown

Disables (or enables) the monitoring session.

**Example:**

```
apic1(config-monitor-access)# no shut
```

**Examples**

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

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

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

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

**Procedure****Step 1** configure

Enters global configuration mode.

**Example:**

```
apic1# configure
```

**Step 2** [no] monitor access filter-group *filtergroup-name*

Creates an access monitoring filter group configuration.

**Example:**

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

**Step 3** `[no] filter srcaddress source-address dstaddress destination-address srcport-from source-from-port srcport-to source-to-port dstport-from destination-from-port dstport-to destination-to-port ipproto IP-protocol`

Configures the filter entries for the filter group, where:

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

**Example:**

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

**Step 4** `exit`

Returns to access monitor filter group configuration mode.

**Example:**

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

**Step 5** `exit`

Exits global configuration mode.

**Example:**

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

---

**Examples**

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

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

```
apic1(config)# exit
```

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

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

### Procedure

- 
- Step 1**     **configure**  
Enters global configuration mode.  
**Example:**  

```
apic1# configure
```
- Step 2**     **[no] monitor access session *session-name***  
Creates an access monitoring session configuration.  
**Example:**  

```
apic1(config)# monitor access session session1
```
- Step 3**     **filter-group *filtergroup-name***  
Associates a filter group.  
**Example:**  

```
apic1(config-monitor-access)# filter-group filtergroup1
```
- Step 4**     **no filter-group**  
Disassociates a filter group, if necessary.  
**Example:**  

```
apic1(config-monitor-access)# no filter-group
```
- Step 5**     **[no] source interface ethernet {[*fx*]/*slot/port* | *port-range*} leaf *node-id***  
Specifies the source interface port or port range.  
**Example:**  

```
apic1(config-monitor-access)# source interface ethernet 1/9 leaf 101
```
- Step 6**     **filter-group *filtergroup-name***  
Associates a filter group to a SPAN source.  
**Example:**  

```
apic1(config-monitor-access-source)# filter-group filtergroup2
```
- Step 7**     **exit**  
Returns to access monitor filter group configuration mode.  
**Example:**  

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

**Step 8**      **no filter-group**  
Disassociates the filter group from a SPAN source, if necessary.

**Example:**

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

**Step 9**      **exit**  
Returns to access monitor filter group configuration mode.

**Example:**

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

**Step 10**     **exit**  
Exits global configuration mode.

**Example:**

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

**Examples**

This example shows how to associate a filter group.

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

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

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

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

**Procedure**

**Step 1**      **configure terminal**  
Enters global configuration mode.

**Example:**

```
apicl# configure terminal
```

- Step 2**      **[no] monitor access session** *session-name*  
Creates an access monitoring session configuration.  
**Example:**  

```
apic1(config)# monitor access session mySession
```
- Step 3**      **[no] description** *text*  
Adds a description for this monitoring session. If the text includes spaces, it must be enclosed in single quotes.  
**Example:**  

```
apic1(config-monitor-access)# description "This is my access ERSPAN session"
```
- Step 4**      **[no] destination tenant** *tenant-name* **application** *application-name* **epg** *epg-name* **destination-ip** *dest-ip-address* **source-ip-prefix** *src-ip-address*  
Specifies the destination interface as a tenant and enters destination configuration mode.  
**Example:**  

```
apic1(config-monitor-access)# destination tenant t1 application appl epg epg1 destination-ip 192.0.20.123 source-ip-prefix 10.0.20.1
```
- Step 5**      **[no] erspan-id** *flow-id*  
Configures the ERSPAN ID for the ERSPAN session. The ERSPAN range is from 1 to 1023.  
**Example:**  

```
apic1(config-monitor-access-dest)# erspan-id 100
```
- Step 6**      **[no] ip dscp** *dscp-code*  
Configures the differentiated services code point (DSCP) value of the packets in the ERSPAN traffic. The range is from 0 to 64.  
**Example:**  

```
apic1(config-monitor-access-dest)# ip dscp 42
```
- Step 7**      **[no] ip ttl** *ttl-value*  
Configures the IP time-to-live (TTL) value for the ERSPAN traffic. The range is from 1 to 255.  
**Example:**  

```
apic1(config-monitor-access-dest)# ip ttl 16
```
- Step 8**      **[no] mtu** *mtu-value*  
Configures the maximum transmit unit (MTU) size for the ERSPAN session. The range is 64 to 9216 bytes.  
**Example:**  

```
apic1(config-monitor-access-dest)# mtu 9216
```
- Step 9**      **exit**  
Returns to monitor access configuration mode.  
**Example:**  

```
apic1(config-monitor-access-dest)#
```
- Step 10**     **[no] source interface ethernet** *{[fex/]slot/port | port-range}* **leaf** *node-id*

Specifies the source interface port or port range.

**Example:**

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

**Step 11** [no] source interface port-channel *port-channel-name-list* leaf *node-id* [*fex fex-id*]

Specifies the source interface port-channel.

**Example:**

```
apicl(config-monitor-access)# source interface port-channel pc1 leaf 101
```

**Step 12** [no] source interface vpc *vpc-name-list* leaf *node-id1* *node-id2* [*fex fex-id1 fex-id2*]

Specifies the source interface vPC.

**Example:**

```
apicl(config-monitor-access)# source interface vpc pc1 leaf 101 102
```

**Step 13** drop enable

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

**Example:**

```
apicl(config-monitor-access-source)# drop enable
```

**Step 14** [no] direction {rx | tx | both}

Specifies direction of traffic to be monitored. The direction can be configured independently for each source port range.

**Example:**

```
apicl(config-monitor-access-source)# direction tx
```

**Step 15** [no] filter tenant *tenant-name* application *application-name* epg *epg-name*

Filters traffic to be monitored. The filter can be configured independently for each source port range.

**Example:**

```
apicl(config-monitor-access-source)# filter tenant t1 application appl epg epg1
```

**Step 16** exit

Returns to access monitor session configuration mode.

**Example:**

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

**Step 17** [no] shutdown

Disables (or enables) the monitoring session.

**Example:**

```
apicl(config-monitor-access)# no shut
```

## Examples

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

```

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

```

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

```

apic1(config-monitor-access)# source interface port-channel pc3 leaf 105

```

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

```

apic1(config-monitor-access)# source interface port-channel vpc3 leaf 105

```

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

```

apic1(config-monitor-access)# source interface eth 101/1/1-2 leaf 105

```

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

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

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

## Procedure

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



- Step 8**      **[no] mtu** *mtu-value*
- Configures the maximum transmit unit (MTU) size for the ERSPAN session. The range is 64 to 9216 bytes.
- Example:**
- ```
apic1(config-monitor-fabric-dest)# mtu 9216
```
- Step 9** **exit**
- Returns to monitor access configuration mode.
- Example:**
- ```
apic1(config-monitor-fabric-dest)#
```
- Step 10**     **[no] source interface ethernet** {*slot/port* | *port-range*} **switch** *node-id*
- Specifies the source interface port or port range.
- Example:**
- ```
apic1(config-monitor-fabric)# source interface eth 1/2 switch 101
```
- Step 11** **drop enable**
- Enables the SPAN on drop feature, which captures all packets that are dropped in the ASIC and sends them to a pre-configured SPAN destination.
- Example:**
- ```
apic1(config-monitor-fabric-source)# drop enable
```
- Step 12**     **[no] direction** {**rx** | **tx** | **both**}
- Specifies direction of traffic to be monitored. The direction can be configured independently for each source port range.
- Example:**
- ```
apic1(config-monitor-fabric-source)# direction tx
```
- Step 13** **[no] filter tenant** *tenant-name* **bd** *bd-name*
- Filters traffic by bridge domain.
- Example:**
- ```
apic1(config-monitor-fabric-source)# filter tenant t1 bd bd1
```
- Step 14**     **[no] filter tenant** *tenant-name* **vrf** *vrf-name*
- Filters traffic by VRF.
- Example:**
- ```
apic1(config-monitor-fabric-source)# filter tenant t1 vrf vrf1
```
- Step 15** **exit**
- Returns to access monitor session configuration mode.
- Example:**
- ```
apic1(config-monitor-fabric-source)# exit
```
- Step 16**     **[no] shutdown**

Disables (or enables) the monitoring session.

**Example:**

```
apicl(config-monitor-fabric)# no shut
```

**Examples**

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

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

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

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

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

### Procedure

- Step 1**     **configure terminal**  
Enters global configuration mode.
- Example:**  
apicl# **configure terminal**
- Step 2**     **[no] monitor tenant** *tenant-name* **session** *session-name*  
Creates a tenant monitoring session configuration.
- Example:**  
apicl(config)# **monitor tenant session** mySession
- Step 3**     **[no] description** *text*

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

**Example:**

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

**Step 4** [no] **destination tenant** *tenant-name* **application** *application-name* **epg** *epg-name* **destination-ip** *dest-ip-address* **source-ip-prefix** *src-ip-address*

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

**Example:**

```
apic1(config-monitor-tenant)# destination tenant t1 application appl epg epg1 destination-ip
192.0.20.123 source-ip-prefix 10.0.20.1
```

**Step 5** [no] **erspan-id** *flow-id*

Configures the ERSPAN ID for the ERSPAN session. The ERSPAN range is from 1 to 1023.

**Example:**

```
apic1(config-monitor-tenant-dest)# erspan-id 100
```

**Step 6** [no] **ip dscp** *dscp-code*

Configures the differentiated services code point (DSCP) value of the packets in the ERSPAN traffic. The range is from 0 to 64.

**Example:**

```
apic1(config-monitor-tenant-dest)# ip dscp 42
```

**Step 7** [no] **ip ttl** *ttl-value*

Configures the IP time-to-live (TTL) value for the ERSPAN traffic. The range is from 1 to 255.

**Example:**

```
apic1(config-monitor-tenant-dest)# ip ttl 16
```

**Step 8** [no] **mtu** *mtu-value*

Configures the maximum transmit unit (MTU) size for the ERSPAN session. The range is 64 to 9216 bytes.

**Example:**

```
apic1(config-monitor-tenant-dest)# mtu 9216
```

**Step 9** **exit**

Returns to monitor access configuration mode.

**Example:**

```
apic1(config-monitor-tenant-dest)#
```

**Step 10** [no] **source application** *application-name* **epg** *epg-name*

Specifies the source interface port or port range.

**Example:**

```
apic1(config-monitor-tenant)# source application app2 epg epg5
```

**Step 11** [no] **direction** {rx | tx | both}

Specifies direction of traffic to be monitored. The direction can be configured independently for each source port range.

**Example:**

```
apicl(config-monitor-tenant-source)# direction tx
```

**Step 12**    **exit**

Returns to access monitor session configuration mode.

**Example:**

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

**Step 13**    **[no] shutdown**

Disables (or enables) the monitoring session.

**Example:**

```
apicl(config-monitor-tenant)# no shut
```

**Examples**

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

```
apicl# configure terminal
apicl(config)# monitor access session mySession
apicl(config-monitor-tenant)# description "This is my tenant ERSPAN session"
apicl(config-monitor-tenant)# destination tenant t1 application app1 epg epg1 destination-ip
192.0.20.123 source-ip-prefix 10.0.20.1
apicl(config-monitor-tenant-dest)# erspan-id 100
apicl(config-monitor-tenant-dest)# ip dscp 42
apicl(config-monitor-tenant-dest)# ip ttl 16
apicl(config-monitor-tenant-dest)# mtu 9216
apicl(config-monitor-tenant-dest)# exit
apicl(config-monitor-tenant)# source application app2 epg epg5
apicl(config-monitor-tenant-source)# direction tx
apicl(config-monitor-tenant-source)# exit
apicl(config-monitor-tenant)# no shut
```

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

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

**Procedure****Step 1**    **configure terminal**

Enters global configuration mode.

**Example:**

```
apicl# configure terminal
```

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

Creates a fabric monitoring session configuration.

**Example:**

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

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

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

**Example:**

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

**Step 4** **source global-drop switch**

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

**Example:**

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

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

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

**Example:**

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

---

**Examples**

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

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

## Configuring SPAN Using the REST API

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

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

## Procedure

Configure a fabric destination group for an ERSPAN destination:

```
POST https://<APIC_IP>/api/node/mo/uni/infra.xml
<spanDestGrp annotation="" descr="" name="Dest2" nameAlias="" ownerKey="" ownerTag="">
 <spanDest annotation="" descr="" name="Dest2" nameAlias="" ownerKey="" ownerTag="">
 <spanRsDestEpg annotation="" dscp="unspecified" finalIp="0.0.0.0" flowId="1"
ip="179.10.10.179"
 mtu="1518"srcIpPrefix="20.20.20.2" tDn="uni/tn-ERSPAN/ap-A1/epg-E1" ttl="64" ver="ver2"

 verEnforced="no"/>
 </spanDest>
</spanDestGrp>
```

## Configuring a Global Drop Source Group Using the REST API

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

## Procedure

Configure a global drop source group:

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

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

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

## Procedure

Configure a leaf port as a SPAN destination:

```
POST https://<APIC_IP>/api/node/mo/uni/infra.xml
<spanDestGrp annotation="" descr="" name="Dest4" nameAlias="" ownerKey="" ownerTag="">
 <spanDest annotation="" descr="" name="Dest4" nameAlias="" ownerKey="" ownerTag="">
 <spanRsDestPathEp annotation="" mtu="1518"
```

```
tDn="topology/pod-1/paths-301/pathep-[eth1/18]"/>
 </spanDest>
</spanDestGrp>
```

## Configuring a SPAN Access Source Group Using the REST API

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

### Procedure

Configure a SPAN access source group:

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

## Configuring a SPAN Fabric Source Group Using the REST API

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

### Procedure

Configure a SPAN fabric source group:

```
POST https://<APIC_IP>/api/node/mo/uni/infra.xml
<spanSrcGrp adminSt="enabled" annotation="" descr="" name="Test-Src2" nameAlias="" ownerKey=""
ownerTag="">
 <spanSrc annotation="" descr="" dir="both" name="Src1" nameAlias="" ownerKey=""
ownerTag="" spanOnDrop="yes">
 <spanRsSrcToPathEp annotation="" tDn="topology/pod-1/paths-301/pathep-[eth1/51]"/>
 </spanSrc>
 <spanSpanLbl annotation="" descr="" name="Dest2" nameAlias="" ownerKey="" ownerTag=""
tag="yellow-green"/>
</spanSrcGrp>
```

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

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

### Procedure

---

Configure an access destination group for an ERSPAN destination.

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

---

## 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.

## acidiag Command

To troubleshoot operations on the Cisco APIC, use the **acidiag** command.



**Caution** This command is not intended for every day operation of ACI. Running all forms of the command can be very disruptive and cause major issues in your network if not used properly. Make sure you understand the full effect on your fabric before running them.

### Cluster Commands

```
acidiag
```

```
acidiag avread
```

```
acidiag fnvread
```

```
acidiag fnvreadex
```

Syntax Description	Option	Function
	<b>avread</b>	<p>Displays APICs within the cluster. The avread output includes:</p> <ul style="list-style-type: none"> <li>• Cluster of—Operational cluster size</li> <li>• out of targeted—The desired cluster size</li> <li>• active= —Indicates whether the APIC is reachable</li> <li>• health= —The overall APIC health summary. Displays services with degraded health scores.</li> <li>• chassisID= —The known chassis IDs for a given APIC.</li> </ul> <p><b>Note</b> Peer chassis IDs can be incorrect for APICs not currently in the cluster.</p>

Option	Function
<b>bootcurr</b>	On the next boot, the APIC system will boot the current APIC image in the Linux partition. This option is not expected to normally be used.
<b>bootother</b>	On the next boot, the APIC system will boot the previous APIC image in the Linux partition. This option is not expected to normally be used.
<b>bond0test</b>	Disruptive test of the APIC connection to the leaf. This is used for internal Cisco testing purposes only and outside of that could cause issues with the APIC connection to the fabric.
<b>fnvread</b>	Displays the address and state of switch nodes registered with the fabric.
<b>fnvreadex</b>	Displays additional information for switch nodes registered with the fabric.
<b>linkflap</b>	Brings down and back up a specified APIC interface.
<b>preservelogs</b>	APIC will archive current logs. During a normal reboot this automatically occurs. This option can be used prior to a hard reboot.
<b>run</b>	Two available options are iptables-list and lldptool. The iptables-list is used to display the Linux iptables, which are controlled by the mgmt Tenant contracts. lldptool is used to display lldp information which is sent or received by the APIC.
<b>rvread</b>	Summarizes the data layer state. The output shows a summary of the data layer state for each service. The shard view shows replicas in ascending order.
<b>acidiag rvread</b> <i>service</i>	Displays the data layer state for a service on all shards across all replicas.  <b>Note</b> For an example, see <a href="#">Examples, on page 119</a>
<b>acidiag rvread</b> <i>service shard</i>	Displays the data layer state for a service on a specific shard across all replicas.  <b>Note</b> For an example, see <a href="#">Examples, on page 119</a>
<b>acidiag rvread</b> <i>service shard replica</i>	Displays the data layer state for a service on a specific shard and replica.  <b>Note</b> For an example, see <a href="#">Examples, on page 119</a>

Option	Function
<b>validateimage</b>	Prior to loading an image into the firmware repository, the image can be validated. Note that this function runs as a normal part of the process of the image being added into the repository.
<b>validateenginxconf</b>	Validates the generated nginx configuration file on APIC to ensure nginx can start with that configuration file. This is meant for debug use, in cases where the nginx webserver is not running on APIC.

### Service IDs

The service IDs listed in the table below are also visible when entering the **man acidiag** command.

*Table 4: Service IDs*

Service	ID
cliD	1
controller	2
eventmgr	3
extXMLApi	4
policyelem	5
policymgr	6
reader	7
ae	8
topomgr	9
observer	10
dbgr	11
observerelem	12
dbgrelem	13
vmmmgr	14
nxosmock	15
bootmgr	16
appliancedirector	17
adrelay	18

<b>Service</b>	<b>ID</b>
ospaagent	19
vleafelem	20
dhcpd	21
scripthandler	22
idmgr	23
ospaelem	24
osh	25
opflexagent	26
opflexelem	27
confelem	28
vtap	29
snmpd	30
opflexp	31
analytics	32
policydist	33
plghandler	34
domainmgr	35
licensemgr	36
N/A	37
platformmgr	38
edmgr	39

**Table 5: Data States**

<b>State</b>	<b>ID</b>
COMATOSE	0
NEWLY_BORN	1
UNKNOWN	2
DATA_LAYER_DIVERGED	11
DATA_LAYER_DEGRADED_LEADERSHIP	12

State	ID
DATA_LAYER_ENTIRELY_DIVERGED	111
DATA_LAYER_PARTIALLY_DIVERGED	112
DATA_LAYER_ENTIRELY_DEGRADED_LEADERSHIP	121
DATA_LAYER_PARTIALLY_DEGRADED_LEADERSHIP	122
FULLY_FIT	255

### System Keywords

```
acidiag [start|stop|restart] [mgmt|xineta]
```

```
acidiag installer -u imageurl -c
```

```
acidiag reboot
```

```
acidiag touch [clean|setup]
```

```
acidiag verifyapic
```

Syntax Description	Option	Function
	<b>-c</b>	Specifies a clean install
	<b>-u</b>	Specifies a URL for the APIC image.
	<i>imageurl</i>	Specifies an APIC image.
	<b>installer</b>	Installs a new image on the APIC, -c for clean install
	<b>mgmt</b>	Specifies all services on the APIC.
	<b>reboot</b>	Reboots the APIC.
	<b>restart</b>	Restarts services on an APIC.
	<b>start</b>	Starts services on an APIC.
	<b>stop</b>	Stops services on an APIC.
	<b>touch [clean   setup]</b>	Resets the APIC configuration. <ul style="list-style-type: none"> <li>• The <b>clean</b> option removes all policy data while retaining the APIC network configuration (such as fabric name, IP address, login)</li> <li>• The <b>setup</b> option removes both policy data and the APIC network configuration.</li> </ul>
	<b>verifyapic</b>	Displays the APIC software version.

Option	Function
<b>xinetd</b>	Specifies xinetd (extended internet daemon) service, which controls the ssh and telnet daemons. Beginning with the 5.3(1) release, telnet is not supported.

### Diagnostic Keywords

acidiag crashsuspecttracker

acidiag dbgtoken

acidiag version

Syntax Description	Option	Function
	<b>crashsuspecttracker</b>	Tracks states of a service or data subset that indicate a crash.
	<b>dbgtoken</b>	Generates a token used to generate a root password. This is to be used as directed while working with the TAC as needed.
	<b>version</b>	Displays the APIC ISO software version.

### Examples

The following examples show how to use the **acidiag** command:

```
apic1# acidiag version 2.2.1o
```

```
apic1# acidiag verifyapic
openssl_check: certificate details
subject= CN=ABC12345678,serialNumber=PID:APIC-SERVER-L1 SN:ABC12345678
issuer= CN=Cisco Manufacturing CA,O=Cisco Systems
notBefore=Sep 28 17:17:42 2016 GMT
notAfter=Sep 28 17:27:42 2026 GMT
openssl_check: passed
ssh_check: passed
all_checks: passed
```

```
apic1# acidiag avread
Local appliance ID=1 ADDRESS=10.0.0.1 TEP ADDRESS=10.0.0.0/16 ROUTABLE IP ADDRESS=0.0.0.0
CHASSIS_ID=1009f750-adab-11e9-a044-8dbd212cd556
Cluster of 7 lm(t):1(2019-08-08T01:02:17.961-07:00) appliances (out of targeted 7
lm(t):7(2019-08-08T03:50:57.240-07:00)) with FABRIC_DOMAIN name=ACI Fabric1 set to
version=apic-4.2(0.235j) lm(t):1(2019-08-17T01:09:16.413-07:00); discoveryMode=PERMISSIVE
lm(t):0(1969-12-31T17:00:00.007-07:00); drrMode=OFF lm(t):0(1969-12-31T17:00:00.007-07:00);
kafkaMode=OFF lm(t):0(1969-12-31T17:00:00.007-07:00)
 appliance id=1 address=10.0.0.1 lm(t):1(2019-08-08T01:02:08.544-07:00) tep
address=10.0.0.0/16 lm(t):1(2019-08-08T01:02:08.544-07:00) routable address=0.0.0.0
lm(t):1(zeroTime) oob address=172.23.96.10/21 lm(t):1(2019-08-08T01:02:18.218-07:00)
version=4.2(0.235j) lm(t):1(2019-08-15T15:22:00.158-07:00)
chassisId=1009f750-adab-11e9-a044-8dbd212cd556 lm(t):1(2019-08-15T15:22:00.158-07:00)
capabilities=0X3EEFFFFFFF--0X2020--0X7F lm(t):1(2019-08-17T01:13:46.997-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-08T01:02:18.228-07:00)
```

```

aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-08T01:02:18.228-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-08T01:02:18.228-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-08T01:02:18.228-07:00)
cntrlSbst=(APPROVED, FCH1748V0SZ) lm(t):1(2019-08-15T15:22:00.158-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=1
lm(t):1(2019-08-08T01:02:08.544-07:00) commissioned=YES lm(t):1(zeroTime) registered=YES
lm(t):1(2019-08-08T01:02:08.544-07:00) standby=NO lm(t):1(2019-08-08T01:02:08.544-07:00)
DRR=NO lm(t):0(zeroTime) apicX=NO lm(t):1(2019-08-08T01:02:08.544-07:00) virtual=NO
lm(t):1(2019-08-08T01:02:08.544-07:00) active=YES(2019-08-08T01:02:08.544-07:00)
health=(applnc:255 lm(t):1(2019-08-17T01:39:26.296-07:00) svc's)
 appliance id=2 address=10.0.0.2 lm(t):7(2019-08-08T03:50:55.470-07:00) tep
address=10.0.0.0/16 lm(t):2(2019-07-23T17:51:38.997-07:00) routable address=0.0.0.0
lm(t):0(zeroTime) oob address=172.23.96.11/21 lm(t):1(2019-08-18T23:14:28.720-07:00)
version=4.2(0.235j) lm(t):2(2019-08-15T15:22:00.300-07:00)
chassisId=694e6a98-adac-11e9-ad79-d1f60e3ee822 lm(t):2(2019-08-15T15:22:00.300-07:00)
capabilities=0X3EEFFFFFFFFF--0X2020--0X2 lm(t):2(2019-08-14T07:55:10.074-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.829-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.829-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.829-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.829-07:00)
cntrlSbst=(APPROVED, FCH1748V0MS) lm(t):2(2019-08-15T15:22:00.300-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=1
lm(t):2(2019-08-08T01:42:03.670-07:00) commissioned=YES lm(t):1(2019-08-08T01:02:17.961-07:00)
 registered=YES lm(t):7(2019-07-24T15:24:25.693-07:00) standby=NO
lm(t):2(2019-08-08T01:42:03.670-07:00) DRR=NO lm(t):1(2019-08-08T01:02:17.961-07:00) apicX=NO
 lm(t):2(2019-08-08T01:42:03.670-07:00) virtual=NO lm(t):0(zeroTime)
active=YES(2019-08-13T17:02:32.983-07:00) health=(applnc:255
lm(t):2(2019-08-17T01:32:51.454-07:00) svc's)
 appliance id=3 address=10.0.0.3 lm(t):7(2019-08-08T03:50:55.470-07:00) tep
address=10.0.0.0/16 lm(t):3(2019-07-23T19:05:56.405-07:00) routable address=0.0.0.0
lm(t):0(zeroTime) oob address=172.23.96.12/21 lm(t):1(2019-08-18T23:14:28.721-07:00)
version=4.2(0.235j) lm(t):3(2019-08-15T15:21:59.893-07:00)
chassisId=1f98b916-adb7-11e9-a6f8-abe00a04e8e6 lm(t):3(2019-08-15T15:21:59.893-07:00)
capabilities=0X3EEFFFFFFFFF--0X2020--0X4 lm(t):3(2019-08-14T07:55:22.256-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
cntrlSbst=(APPROVED, FCH1930V1X6) lm(t):3(2019-08-15T15:21:59.893-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=1
lm(t):3(2019-08-08T02:15:20.560-07:00) commissioned=YES lm(t):2(2019-08-08T01:42:15.337-07:00)
 registered=YES lm(t):7(2019-07-24T15:24:25.693-07:00) standby=NO
lm(t):3(2019-08-08T02:15:20.560-07:00) DRR=NO lm(t):2(2019-08-08T01:42:15.337-07:00) apicX=NO
 lm(t):3(2019-08-08T02:15:20.560-07:00) virtual=NO lm(t):0(zeroTime)
active=YES(2019-08-13T17:02:33.182-07:00) health=(applnc:255
lm(t):3(2019-08-15T16:08:46.119-07:00) svc's)
 appliance id=4 address=10.0.0.4 lm(t):7(2019-08-08T03:50:55.470-07:00) tep
address=10.0.0.0/16 lm(t):4(2019-07-23T17:46:15.545-07:00) routable address=0.0.0.0
lm(t):0(zeroTime) oob address=172.23.97.231/21 lm(t):1(2019-08-18T23:14:28.717-07:00)
version=4.2(0.235j) lm(t):4(2019-08-15T15:22:00.669-07:00)
chassisId=3a7f38aa-adac-11e9-8869-a9e520cdc042 lm(t):4(2019-08-15T15:22:00.669-07:00)
capabilities=0X3EEFFFFFFFFF--0X2020--0X8 lm(t):4(2019-08-14T07:54:59.490-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.825-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.825-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.825-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.825-07:00)
cntrlSbst=(APPROVED, FCH1902V1WW) lm(t):4(2019-08-15T15:22:00.669-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=1
lm(t):4(2019-08-08T02:40:09.610-07:00) commissioned=YES lm(t):3(2019-08-08T02:15:32.613-07:00)
 registered=YES lm(t):7(2019-07-24T15:24:25.693-07:00) standby=NO
lm(t):4(2019-08-08T02:40:09.610-07:00) DRR=NO lm(t):3(2019-08-08T02:15:32.613-07:00) apicX=NO
 lm(t):4(2019-08-08T02:40:09.610-07:00) virtual=NO lm(t):0(zeroTime)
active=YES(2019-08-15T15:21:59.914-07:00) health=(applnc:255
lm(t):4(2019-08-17T01:39:26.477-07:00) svc's)

```



```

 appliance id=5 address=10.0.0.5 lm(t):7(2019-08-08T03:50:55.470-07:00) tep
address=10.0.0.0/16 lm(t):5(2019-07-23T19:05:11.089-07:00) routable address=0.0.0.0
lm(t):0(zeroTime) oob address=172.23.97.232/21 lm(t):1(2019-08-18T23:14:28.723-07:00)
version=4.2(0.235j) lm(t):5(2019-08-15T15:22:00.248-07:00)
chassisId=35428666-adb7-11e9-a315-1d7671b518b3 lm(t):5(2019-08-15T15:22:00.248-07:00)
capabilities=0X3EEEEEEEEEE--0X2020--0X10 lm(t):5(2019-08-14T07:55:19.573-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.854-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.854-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.854-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.854-07:00)
cntrlSbst=(APPROVED, FCH1902V1EG) lm(t):5(2019-08-15T15:22:00.248-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=1
lm(t):5(2019-08-08T03:03:50.338-07:00) commissioned=YES lm(t):4(2019-08-08T02:40:15.939-07:00)
registered=YES lm(t):7(2019-07-24T15:24:25.693-07:00) standby=NO
lm(t):5(2019-08-08T03:03:50.338-07:00) DRR=NO lm(t):4(2019-08-08T02:40:15.939-07:00) apicX=NO
lm(t):5(2019-08-08T03:03:50.338-07:00) virtual=NO lm(t):0(zeroTime)
active=YES(2019-08-15T15:21:59.756-07:00) health=(applnc:255
lm(t):5(2019-08-17T01:32:43.730-07:00) svc's)
 appliance id=6 address=10.0.0.6 lm(t):7(2019-08-08T03:50:55.470-07:00) tep
address=10.0.0.0/16 lm(t):6(2019-07-23T19:39:41.972-07:00) routable address=0.0.0.0
lm(t):0(zeroTime) oob address=172.31.170.230/21 lm(t):1(2019-08-18T23:14:28.727-07:00)
version=4.2(0.235j) lm(t):6(2019-08-15T15:22:00.562-07:00)
chassisId=066c943a-adbc-11e9-bbed-257398025731 lm(t):6(2019-08-15T15:22:00.562-07:00)
capabilities=0X3EEEEEEEEEE--0X2020--0X20 lm(t):6(2019-08-14T07:55:20.053-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.820-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.821-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.821-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.821-07:00)
cntrlSbst=(APPROVED, WZP22350JFT) lm(t):6(2019-08-15T15:22:00.562-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=9
lm(t):6(2019-08-08T03:28:11.246-07:00) commissioned=YES lm(t):5(2019-08-08T03:03:57.387-07:00)
registered=YES lm(t):7(2019-07-24T15:24:25.693-07:00) standby=NO
lm(t):6(2019-08-08T03:28:11.246-07:00) DRR=NO lm(t):5(2019-08-08T03:03:57.387-07:00) apicX=NO
lm(t):6(2019-08-08T03:28:11.246-07:00) virtual=NO lm(t):0(zeroTime)
active=YES(2019-08-13T17:30:37.663-07:00) health=(applnc:255
lm(t):6(2019-08-15T15:57:05.128-07:00) svc's)
 appliance id=7 address=10.0.0.7 lm(t):7(2019-08-08T03:50:48.149-07:00) tep
address=10.0.0.0/16 lm(t):7(2019-07-24T15:24:19.988-07:00) routable address=0.0.0.0
lm(t):0(zeroTime) oob address=172.31.172.157/21 lm(t):1(2019-08-18T23:14:28.722-07:00)
version=4.2(0.235j) lm(t):7(2019-08-15T15:22:00.539-07:00)
chassisId=859be4ae-ae61-11e9-9840-7d9d67698989 lm(t):7(2019-08-15T15:22:00.539-07:00)
capabilities=0X3EEEEEEEEEE--0X2020--0X40 lm(t):7(2019-08-14T07:55:23.872-07:00)
rK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
aK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
oobrK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
oobaK=(stable,present,0X206173722D687373) lm(t):1(2019-08-15T15:22:00.824-07:00)
cntrlSbst=(APPROVED, FCH2051V116) lm(t):7(2019-08-15T15:22:00.539-07:00) (targetMbSn=
lm(t):0(zeroTime), failoverStatus=0 lm(t):0(zeroTime)) podId=10
lm(t):7(2019-08-08T03:50:48.149-07:00) commissioned=YES lm(t):6(2019-08-08T03:28:16.727-07:00)
registered=YES lm(t):6(2019-07-24T15:27:25.518-07:00) standby=NO
lm(t):7(2019-08-08T03:50:48.149-07:00) DRR=NO lm(t):6(2019-08-08T03:28:16.727-07:00) apicX=NO
lm(t):7(2019-08-08T03:50:48.149-07:00) virtual=NO lm(t):0(zeroTime)
active=YES(2019-08-13T17:30:45.488-07:00) health=(applnc:255
lm(t):7(2019-08-17T01:39:26.549-07:00) svc's)

clusterTime=<diff=2817 common=2019-08-19T15:33:55.929-07:00
local=2019-08-19T15:33:53.112-07:00 pF=<displForm=0 offsSt=0 offsVlu=-25200
lm(t):7(2019-08-08T03:50:55.925-07:00)>>

```

```
apic1# acidiag rvread 6 3 1
```

```
(6,3,1) st:6 lm(t):3(2014-10-16T08:48:20.238+00:00) le: reSt:LEADER voGr:0 cuTerm:0x19
lCoTe:0x18
```

```
lCoIn:0x18000000000001b2a veFiSt:0x31 veFiEn:0x31 lm(t):3(2014-10-16T08:48:20.120+00:00)
```

```

 lastUpdt 2014-10-16T09:07:00.214+00:00

clusterTime=<diff=65247252 common=2014-10-16T09:07:01.837+00:00
local=2014-10-15T14:59:34.585+00:00
 pF=<displForm=0 offsSt=0 offsVlu=0 lm(t):3(2014-10-16T04:50:08.714+00:00)>>

apicl# acidiag rvread 6 3
(6,3,1) st:6 lm(t):3(2014-10-16T08:48:20.238+00:00) le: reSt:LEADER voGr:0 cuTerm:0x19
lCoTe:0x18
 lCoIn:0x1800000000001b2a veFiSt:0x31 veFiEn:0x31 lm(t):3(2014-10-16T08:48:20.120+00:00)

 lastUpdt 2014-10-16T09:08:30.240+00:00
(6,3,2) st:6 lm(t):1(2014-10-16T08:47:25.323+00:00) le: reSt:FOLLOWER voGr:0 cuTerm:0x19
lCoTe:0x18
 lCoIn:0x1800000000001b2a veFiSt:0x49 veFiEn:0x49 lm(t):1(2014-10-16T08:48:20.384+00:00)
 lp: clSt:2
 lm(t):1(2014-10-16T08:47:03.286+00:00) dbSt:2 lm(t):1(2014-10-16T08:47:02.143+00:00)
 stMmt:1
 lm(t):(zeroTime) dbCrTs:2014-10-16T08:47:02.143+00:00 lastUpdt
2014-10-16T08:48:20.384+00:00
(6,3,3) st:6 lm(t):2(2014-10-16T08:47:13.576+00:00) le: reSt:FOLLOWER voGr:0 cuTerm:0x19
lCoTe:0x18
 lCoIn:0x1800000000001b2a veFiSt:0x43 veFiEn:0x43 lm(t):2(2014-10-16T08:48:20.376+00:00)

 lastUpdt 2014-10-16T09:08:30.240+00:00

clusterTime=<diff=65247251 common=2014-10-16T09:08:30.445+00:00
local=2014-10-15T15:01:03.194+00:00
 pF=<displForm=0 offsSt=0 offsVlu=0 lm(t):3(2014-10-16T04:50:08.714+00:00)>>

```



## CHAPTER 6

# Provisioning Core ACI Fabric Services

---

This chapter contains the following sections:

- [Link Level Policies](#), on page 123
- [Link Flap Policies](#), on page 124
- [Time Synchronization and NTP](#), on page 125
- [Configuring a DHCP Relay Policy](#), on page 131
- [Configuring a DNS Service Policy](#), on page 139
- [Configuring Custom Certificates](#), on page 143
- [Provisioning Fabric Wide System Settings](#), on page 146
- [Provisioning Global Fabric Access Policies](#), on page 171
- [Per Port Policies](#), on page 175
- [Creating a Mis-cabling Protocol Interface Policy Using the GUI \(Optional\)](#), on page 177

## Link Level Policies

You can configure link level policies, which are a type of access policy. A link level policy includes the physical layer (Layer 1) interface configurations, such as auto-negotiation, port speed, and link debounce.

## Electromagnetic Interference Retrain

In the 5.2(4) and later releases, the electromagnetic interference (EMI) retrain feature filters any noise on a link due to electromagnetic interference, and retrains the link to avoid a link flap. Enable EMI retrain if your data center environment has a lot of EMI noise.

You can enable EMI retrain when you configure a link level policy by choosing **enable** for the **EMI Retrain** property. This feature is supported only with the Cisco N9K-C93108TC-EX and N9K-C93108TC-FX leaf switches using copper cables.

## Configuring a Link Level Policy Using the GUI

### Procedure

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

**Step 2** In the **Navigation** pane, choose **Policies > Interface > Link Level**.

**Step 3** Right-click **Link Level** and choose **Create Link Level Policy**.

**Step 4** In the **Create Link Level Policy** dialog, fill out the fields as appropriate for your desired configuration.

For **Speed**, we recommend that you choose **inherit**, which is the default value. With this value, the Cisco APIC determines the speed based on the transceiver that is inserted into a switch.

See the tooltips for more information about the fields.

**Step 5** Click **Submit**.

## Port Bring-up Delay

Beginning with the 4.2(5) release, when you configure a link level policy, you can set the **Port bring-up delay (milliseconds)** parameter, which specifies a time in milliseconds that the decision feedback equalizer (DFE) tuning is delayed when a port is coming up. The delay is used to help avoid CRC errors during link bringup when using some third-party adapters. You should set the delay only as required; in most cases, you do not need to set a delay.



**Note** The **Port bring-up delay (milliseconds)** parameter is not honored on fabric extender (FEX) ports.

## Link Debounce Interval

A link level policy applies to the fabric interfaces (leaf switch uplink ports and spine switch ports) at the link level and you use this policy where there might be link noise or transient flaps. In this policy, you can configure the link debounce interval, which specifies a wait time for the fabric interfaces when there is a flap in the link. After the interval passes, Cisco Application Centric Infrastructure (ACI) checks the link status again to re-confirm the event. If the link is okay at this time, then the interfaces will remain up.

Set the interval in a link level policy using the **Link debounce interval (msec)** field in the Cisco Application Policy Infrastructure Controller (APIC) GUI or the **[no] link debounce time <time>** CLI command. The default debounce interval is 0 ms. We recommend a value of 100 ms, but you should choose a value that is appropriate to your fabric.

## Link Flap Policies

Link flap is a situation in which a physical interface on a switch continually goes up and down over a period of time. The cause is usually related to a bad, unsupported, or non-standard cable or Small Form-Factor Pluggable (SFP), or is related to other link synchronization issues, and the cause can be intermittent or permanent.

A link flap policy specifies when to disable a switch port due to link flapping errors. In a link flap policy, you specify maximum number of times that a port of a switch can flap within a specified time span. If the port flaps more than the specified number of times in the specified time span, the port is given the "error-disable" state. The port remains in this state until you perform a manual flap on the port using the Cisco Application Policy Infrastructure Controller (APIC) to disable and enable the port.



**Note** A link flap policy is not honored on fabric extender (FEX) host interface (HIF) ports nor on leaf switch models without -EX, -FX, -FX2, -GX, or later designations in the product ID.

## Configuring a Link Flap Policy Using the GUI

The following procedure configures a link flap policy using the GUI, which you can then attach to any leaf or spine node interface policy to deploy the link flap policy on the node's access ports.

### Procedure

- Step 1** On the menu bar, choose **Fabric > Access Policies**.
- Step 2** In the **Navigation** pane, choose **Policies > Interface > Link Flap**.
- Step 3** Right-click **Link Flap** and choose **Create Link Flap Policy**.
- Step 4** In the **Create Link Level Policy** dialog, fill out the fields as appropriate for your desired configuration.  
See the tooltips and online help for more information about the fields.
- Step 5** Click **Submit**.

## Time Synchronization and NTP

Within the Cisco Application Centric Infrastructure (ACI) fabric, time synchronization is a crucial capability upon which many of the monitoring, operational, and troubleshooting tasks depend. Clock synchronization is important for proper analysis of traffic flows as well as for correlating debug and fault time stamps across multiple fabric nodes.

An offset present on one or more devices can hamper the ability to properly diagnose and resolve many common operational issues. In addition, clock synchronization allows for the full utilization of the atomic counter capability that is built into the ACI upon which the application health scores depend. Nonexistent or improper configuration of time synchronization does not necessarily trigger a fault or a low health score. You should configure time synchronization before deploying a full fabric or applications so as to enable proper usage of these features. The most widely adapted method for synchronizing a device clock is to use Network Time Protocol (NTP).

Prior to configuring NTP, consider what management IP address scheme is in place within the ACI fabric. There are two options for configuring management of all ACI nodes and Application Policy Infrastructure Controllers (APICs), in-band management and/or out-of-band management. Depending upon which management option is chosen for the fabric, configuration of NTP will vary. Another consideration in deploying time synchronization is where the time source is located. The reliability of the source must be carefully considered when determining if you will use a private internal clock or an external public clock.

## In-Band Management NTP



**Note** See the Adding Management Access section in this guide for information about in-band management access.

- **In-Band Management NTP**—When an ACI fabric is deployed with in-band management, consider the reachability of the NTP server from within the ACI in-band management network. In-band IP addressing used within the ACI fabric is not reachable from anywhere outside the fabric. To leverage an NTP server external to the fabric with in-band management, construct a policy to enable this communication..

## NTP over IPv6

NTP over IPv6 addresses is supported in hostnames and peer addresses. The `gai.conf` can also be set up to prefer the IPv6 address of a provider or a peer over an IPv4 address. The user can provide a hostname that can be resolved by providing an IP address (both IPv4 or IPv6, depending on the installation or preference).

## Configuring NTP Using the GUI



**Note** There is a risk of hostname resolution failure for hostname based NTP servers if the DNS server used is configured to be reachable over in-band or out-of-band connectivity. If you use a hostname, ensure that the DNS service policy to connect with the DNS providers is configured. Also ensure that the appropriate DNS label is configured for the in-band or out-of-band VRF instances of the management EPG that you chose when you configured the DNS profile policy.

### Procedure

- 
- Step 1** On the menu bar, choose **Fabric > Fabric Policies**.
- Step 2** In the **Navigation** pane, choose **Policies > Pod > Date and Time**.
- Step 3** In the **Work** pane, choose **Actions > Create Date and Time Policy**.
- Step 4** In the **Create Date and Time Policy** dialog box, perform the following actions:
- Enter a name for the policy to distinguish between the different NTP configurations in your environment..
  - Click **enabled** for the **Authentication State** field and expand the **NTP Client Authentication Keys** table and enter the key information. Click **Update** and **Next**.
  - Click the + sign to specify the NTP server information (provider) to be used.
  - In the **Create Providers** dialog box, enter all relevant information, including the following fields: **Name**, **Description**, **Minimum Polling Intervals**, and **Maximum Polling Intervals**.
    - If you are creating multiple providers, check the **Preferred** check box for the most reliable NTP source.
    - In the **Management EPG** drop-down list, if the NTP server is reachable by all nodes on the fabric through out-of-band management, choose Out-of-Band. If you have deployed in-band management, see the details about In-Band Management NTP. Click **OK**.

Repeat the steps for each provider that you want to create.

- Step 5** In the **Navigation** pane, choose **Pod Policies > Policy Groups**.
- Step 6** In the **Work** pane, choose **Actions > Create Pod Policy Group**.
- Step 7** In the **Create Pod Policy Group** dialog box, perform the following actions:
- Enter a name for the policy group.
  - In the **Date Time Policy** field, from the drop-down list, choose the NTP policy that you created earlier. Click **Submit**.  
The pod policy group is created. Alternatively, you can use the default pod policy group.
- Step 8** In the **Navigation** pane, choose **Pod Policies > Profiles**.
- Step 9** In the **Work** pane, double-click the desired pod selector name.
- Step 10** In the Properties area, from the **Fabric Policy Group** drop-down list, choose the pod policy group you created. Click **Submit**.

## Configuring NTP Using the REST API



**Note** There is a risk of hostname resolution failure for hostname based NTP servers if the DNS server used is configured to be reachable over in-band or out-of-band connectivity. If you use a hostname, ensure that the DNS service policy to connect with the DNS providers is configured. Also ensure that the appropriate DNS label is configured for the in-band or out-of-band VRF instances of the management EPG that you chose when you configured the DNS profile policy.

### Procedure

- Step 1** Configure NTP.

**Example:**

POST url: `https://APIC-IP/api/node/mo/uni/fabric/time-test.xml`

```
<imdata totalCount="1">
 <datetimePol adminSt="enabled" authSt="disabled" descr="" dn="uni/fabric/time-CiscoNTPPol"
 name="CiscoNTPPol" ownerKey="" ownerTag="">
 <datetimeNtpProv descr="" keyId="0" maxPoll="6" minPoll="4" name="10.10.10.11"
 preferred="yes">
 <datetimeRsNtpProvToEpg tDn="uni/tn-mgmt/mgmtmp-default/inb-default"/>
 </datetimeNtpProv>
 </datetimePol>
</imdata>
```

- Step 2** Add the default Date Time Policy to the pod policy group.

**Example:**

POST url: `https://APIC-IP/api/node/mo/uni/fabric/funcprof/podgrp-calol/rsTimePol.xml`

```
POST payload: <imdata totalCount="1">
<fabricRsTimePol tnDatetimePolName="CiscoNTPPol">
```

```
</fabricRsTimePol>
</imdata>
```

**Step 3** Add the pod policy group to the default pod profile.

**Example:**

```
POST url:
https://APIC-IP/api/node/mo/uni/fabric/podprof-default/pods-default-ty-ALL/rsPodPGrp.xml

payload: <imdata totalCount="1">
<fabricRsPodPGrp tDn="uni/fabric/funcprof/podpgrp-calol" status="created">
</fabricRsPodPGrp>
</imdata>
```

## Verifying NTP Operation Using the GUI

### Procedure

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

**Step 2** In the **Navigation** pane, choose **Pod Policies > Policies > Date and Time > ntp\_policy > server\_name**.

The *ntp\_policy* is the previously created policy. An IPv6 address is supported in the Host Name/IP address field. If you enter a hostname and it has an IPv6 address set, you must implement the priority of IPv6 address over IPv4 address.

**Step 3** In the **Work** pane, verify the details of the server.

## NTP Server

The NTP server enables client switches to also act as NTP servers to provide NTP time information to downstream clients. When the NTP server is enabled, the NTP daemon on the switch responds with time information to all unicast (IPv4/IPv6) requests from NTP clients. NTP server implementation is compliant to NTP RFCv3. As per the NTP RFC, the server will not maintain any state related to the clients.

- The NTP server enables IP addresses in all tenant VRFs and the in-band/out-of-band management VRFs to serve NTP clients.
- The NTP server responds to incoming NTP requests on both Management VRFs or tenant VRFs, and responds back using the same VRF.
- The NTP server supports both IPv4/IPv6.
- Switches can sync as an IPv4 client and serve as an IPv6 server, and vice versa.
- Switches can sync as an NTP client using either the out-of-band management or in-band management VRF and serve NTP clients from either management VRF or tenant VRF.
- No additional contracts or IP table configurations are required.



- If the switch is synced to the upstream server, then the server will send time info with the stratum number, and increment to its system peer's stratum.
- If the switch clock is undisciplined (not synced to the upstream server), then the server will send time information with stratum 16. Clients will not be able to sync to this server.

By default, NTP server functionality is disabled. It needs to be enabled explicitly by the configuration policy.



**Note** Clients can use the in-band, out-of-band IP address of the leaf switch as the NTP server IP address. Clients can also use the bridge domain SVI of the EPG of which they are part or any L3Out IP address as the NTP server IP address for clients outside of the fabric.

Fabric switches should not sync to other switches of the same fabric. The fabric switches should always sync to external NTP servers.

## Enabling the NTP Server Using the GUI

This section explains how to enable an NTP server when configuring NTP in the APIC GUI.

### Procedure

- Step 1** On the menu bar, choose **FABRIC > Fabric Policies**.
- Step 2** In the **Navigation** pane, choose **Pod Policies > Policies**.  
The **Date and Time** option appears in the **Navigation** pane.
- Step 3** From the **Navigation** pane, right-click on **Date and Time** and choose **Create Date and Time Policy**.  
The **Create Date and Time Policy** dialog appears in the **Work** pane.
- Step 4** In the **Create Date and Time Policy** dialog box, perform the following actions:
  - a) Enter a name for the policy to distinguish between the different NTP configurations in your environment.
  - b) For the **Server State** option, click **enabled**.

**Server State** enables switches to act as NTP servers to provide NTP time information to downstream clients.

**Note**

To support the server functionality, it is always recommended to have a peer setup for the server. This enables the server to have a consistent time to provide to the clients.

When **Server State** is enabled:

- The NTP server sends time info with a stratum number, an increment to the system peer's stratum number, to switches that are synced to the upstream server.
- The server sends time info with stratum 16 if the switch clock is not synced to the upstream server. Clients are not able to sync to this server.

**Note**

To support the server functionality, it is always recommended to have a peer setup for the server. The peer setup allows for a consistent time to provide to the clients.

- c) For the **Master Mode** option, click **enabled**.

**Master Mode** enables the designated NTP server to provide undisciplined local clock time to downstream clients with a configured stratum number. For example, a leaf switch that is acting as the NTP server can provide undisciplined local clock time to leaf switches acting as clients.

**Note**

- **Master Mode** is only applicable when the server clock is undisciplined.
- The default master mode **Stratum Value** is 8.

- d) For the **Stratum Value** field, specify the stratum level from which NTP clients will get their time synchronized. The range is from 1 to 14.
- e) Click **Next**.
- f) Click the + sign to specify the NTP server information (provider) to be used.
- g) In the **Create Providers** dialog box, enter all relevant information, including the following fields: **Name**, **Description**, **Minimum Polling Intervals**, and **Maximum Polling Intervals**.
- If you are creating multiple providers, check the **Preferred** check box for the most reliable NTP source.
  - In the Management EPG drop-down list, if the NTP server is reachable by all nodes on the fabric through out-of-band management, choose Out-of-Band. If you have deployed in-band management, see the details about In-Band Management NTP. Click **OK**.

Repeat the steps for each provider that you want to create.

**Step 5** In the **Navigation** pane, choose **Pod Policies** then right-click on **Policy Groups**.

The **Create Pod Policy Group** dialog appears.

**Step 6** In the **Work** pane, choose **Actions > Create Pod Policy Group**.

**Step 7** In the **Create Pod Policy Group** dialog box, perform the following actions:

- a) Enter a name for the policy group.
- b) In the **Date Time Policy** field, from the drop down list, choose the NTP policy that you created earlier. Click **Submit**.

The pod policy group is created. Alternatively, you can use the default pod policy group.

**Step 8** In the **Navigation** pane, choose **Pod Policies > Profiles**.

**Step 9** In the **Work** pane, double-click the desired pod selector name.

**Step 10** In the Properties area, from the **Fabric Policy Group** drop down list, choose the pod policy group you created.

**Step 11** Click **Submit**.

---

## Configuring the Datetime Format Using the GUI

This section demonstrates how to configure the datetime format using the Cisco APIC GUI.

## Procedure

- 
- Step 1** On the menu bar, click **System > System Settings** .
- Step 2** In the Navigation pane, click **Date and Time**.
- Step 3** In the Work pane, choose from the following options:
- **Display Format**—Click **local** to display the date and time in local time, or click **utc** to display the date and time in UTC. The default is **local**.
  - **Time Zone**—Click the drop-down arrow to choose the time zone for your domain. The default is **Coordinated Universal Time**.
  - **Offset State**—Click **enable** or **disable**. When enabled, the difference between the local time and the reference time is displayed. The default is **enable**.
- 

# Configuring a DHCP Relay Policy

A DHCP relay policy may be used when the DHCP client and server are in different subnets.

When a Cisco Application Centric Infrastructure (ACI) fabric acts as a DHCP relay, it inserts the DHCP Option 82 (the DHCP Relay Agent Information Option) in DHCP requests that it proxies on behalf of clients. If a response (DHCP offer) comes back from a DHCP server without Option 82, it is silently dropped by the fabric. Therefore, when the Cisco ACI fabric acts as a DHCP relay, DHCP servers providing IP addresses to compute nodes attached to the Cisco ACI fabric must support Option 82.

## Option 79: Client Link-Layer Address Option in DHCPv6

Beginning with Cisco Application Policy Infrastructure Controller (APIC) release 5.2(4), you can configure bridge domains configured as DHCP relay agents to include DHCPv6 Option 79. When Option 79 is enabled, the leaf switch with the bridge domain configured as the relay agent will include the client's link-layer address through option 79 of the DHCPv6 relay packet.

When option 79 is selected, the payload of the DHCP packet contains the client mac address (client link layer address). Option 79 contains the actual link-layer address of the device. The relay message uses the Ethernet Source MAC address in the actual DHCP packets coming from the client and prefixes it with 00:01 to indicate an ethernet source and then copies these 8 bytes (client mac address) in to Option 79.

For more details about Client Link-Layer Address Option in DHCPv6, see *RFC 6939*.

## Benefit of using Option 79

In dual-stack scenarios (support for IPv6 and IPv4), when you need to associate DHCPv4 and DHCPv6 messages with the same client interface, option 79 carries the client MAC address in the DHCPv6 relay packets, complying with RFC standards.

## About the DHCP Server Preference Field




**Note** Following are definitions for several terms used in this section:

- **Client VRF:** The VRF instance where the host initiating a DHCP request is located.
- **Server VRF:** The VRF instance where either the DHCP server is located or the VRF instance that provides the path to reach the DHCP server (for example, via an L3Out).
- **Client EPG:** The EPG where the host initiating a DHCP request is located.
- **Server EPG:** The EPG where the DHCP server is connected (or alternatively the external EPG if the DHCP server is outside of the ACI fabric).

Cisco APIC release 5.2(4) adds support for the `use-vrf` option when configuring a DHCP relay provider. This feature is used when the DHCP provider EPG (for example, the EPG where the DHCP server is connected to) or Layer 3 external network through which the DHCP server is reachable is in a different VRF instance than the bridge domain where the host initiating the DHCP request is located (the bridge domain that is referencing the DHCP policy as a DHCP relay label). This feature is equivalent to the DHCP relay `use-vrf` option available in NX-OS. When the `use-vrf` option is enabled for a DHCP relay provider, the leaf switch, where the DHCP client is located, will route the DHCP relay packet via the VRF instance of the configured DHCP provider EPG (or the configured L3Out that provides reachability to the DHCP server) instead of the VRF instance of the DHCP client.

Cisco APIC releases prior to release 5.2(4) support specifying a DHCP relay provider (server) in an EPG or Layer 3 external network of a different VRF instance than the VRF instance where the DHCP client is. This inter-VRF relay policy relies on inter-VRF contracts and route leaking of the DHCP server network from the VRF instance that has reachability to the DHCP server (also known as the **server VRF**) into the VRF instance where the DHCP client is (also known as the **client VRF**). The DHCP relay packet is routed from the client VRF instance and uses the inter-VRF route leaking to reach the DHCP server from the server VRF instance. In some scenarios, the DHCP relay packet might bypass the server VRF instance if the DHCP server network is also reachable from the client VRF instance (for example, if there is a local L3Out in the client VRF instance that can also reach the DHCP server network). When a DHCP relay policy provider is configured to use a Layer 3 external network different from the one of the client VRF instance, the source IP address of the DHCP relay packets is selected from the L3Out of the server VRF instance (also known as the **provider L3Out**). When these DHCP relay packets are routed from the L3Out of the client VRF instance instead of the L3Out of the server VRF instance (which can happen if the L3Out of the client VRF instance also has a route to the DHCP server), the DHCP server response will be sent back to the L3Out in the server VRF instance because the IP address in the DHCP relay packet is set to the IP address of the L3Out of the server VRF instance. This can result in asymmetric forwarding of the DHCP relay packets and may be dropped by stateful devices, such as a firewall.

The following figure provides an example of this scenario.



In this example scenario, the external DHCP server network is reachable in the ACI fabric via both the client and the server VRF instance. The DHCP relay packet is routed from the client VRF instance and sent via the client VRF instance L3Out. The source IP address in the DHCP relay packet is selected from the server VRF instance L3Out per the DHCP relay policy. The DHCP relay response from the server will be routed to the DHCP server L3Out, resulting in an asymmetric flow.

To resolve this issue, the **Use Server VRF** option in the **DHCP Server Preference** field is available beginning with release 5.2(4). With the **Use Server VRF** option enabled, the DHCP relay packet will always be routed from the server VRF instance. This option also removes the requirement for inter-VRF contracts and route leaking.

Based on the option that you select in the **DHCP Server Preference** field, the leaf switch determines whether to route the DHCP relay packets from the client VRF instance or the server VRF instance:

- **None:** This is the default option, which reflects the behavior prior to release 5.2(4). By choosing the **None** option, the switch will always route the DHCP relay packet from the client VRF instance. If used for inter-VRF DHCP relay, a shared services contract is required to leak the server VRF instance network into the client VRF instance.
- **Use Server VRF:** This option reflects new behavior introduced in release 5.2(4). By choosing the **Use Server VRF** option, the switch routes the DHCP relay packets from the server VRF instance, regardless of whether there is a contract or there is no contract between the EPG where the DHCP client is and the

EPG where the DHCP server is (or the Layer 3 external of the L3Out through which the DHCP server is reachable).

For inter-VRF configurations, when you choose the **Use Server VRF** option in the **DHCP Server Preference** field, the server subnet route is programmed in the server VRF instance on the client leaf switch for route lookup. The DHCP process on the client leaf switch then sends the DHCP relay packets via the server VRF instance. Because of this, the server VRF instance must be deployed with at least one IP address on all leaf switches where the client bridge domains are deployed.

## Guidelines and Limitations for a DHCP Relay Policy

- A DHCP relay policy created under the infra or common tenant is not available to other tenants when configuring DHCP relay in a bridge domain. For inter-tenant DHCP relay communications, create a global DHCP relay policy, as described in [Create a Global DHCP Relay Policy, on page 172](#).
- The DHCP relay IP address will always be set to the primary SVI IP address.
- You cannot use a DHCP relay policy if the clients and servers are connected through an L3Out EPG. Either the clients or servers (or both) must be in a regular (non-L3Out) EPG.
- When a DHCP relay provider EPG is an L3Out External EPG and is configured in a different VRF instance than the client bridge domain, a DHCP relay packet will be sent sourced from each interface configured under the L3Out. This includes the node router-IDs addresses and all L3Out interface IP addresses.

## Configuring a DHCP Server Policy for the APIC Infrastructure Using the GUI

This procedure deploys a DHCP relay policy for an endpoint group (EPG).

Observe the following guidelines and restrictions:

- The port and the encapsulation used by the application Endpoint Group must belong to a physical or VM Manager (VMM) domain. If no such association with a domain is established, the APIC continues to deploy the EPG but raises a fault.
- Cisco APIC supports DHCP relay for both IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6. DHCPv6 relay will occur only if IPv6 is enabled on the fabric interface and one or more DHCPv6 relay servers are configured.
- Cisco APIC supports DHCP relay for only the primary IP address pool.
- The following guidelines and restrictions apply for the **DHCP Server Preference** field introduced in release 5.2(4).
  - When a DHCP relay is configured for an L3Out case (for example, when the DHCP server is behind an L3Out and a DHCP relay policy is configured with the **Use Server VRF** option in the **DHCP Server Preference** field), then you must deploy an EPG/bridge domain/bridge domain subnet to the leaf switches where the client bridge domains are deployed if the server VRF has no interfaces already present.
  - When a DHCP relay policy is configured with the **Use Server VRF** option in the **DHCP Server Preference** field, for a DHCP server that is behind an EPG, then both the IPv4 and IPv6 routes and the server bridge domain SVI is created on the client leaf switch.

- The **Use Server VRF** option is not supported with intersite DHCP traffic.
- The following restrictions apply for Option 79:
  - Option 79 is supported on DHCPv6 only.
  - Option 79 is not supported on *infra* tenant.

### Before you begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

## Procedure

- 
- Step 1** On the menu bar, choose **Tenants** > *tenant\_name*.
- Step 2** In the **Navigation** pane, under **Tenant** *tenant\_name*, expand **Policies** > **Protocol** > **DHCP**.
- Step 3** Right-click **Relay Policies** and click **Create DHCP Relay Policy**.
- Step 4** In the **Create DHCP Relay Policy** dialog box, perform the following actions:
- a) In the **Name** field, enter the DHCP relay profile name (DhcpRelayP).  
This name can be up to 64 alphanumeric characters.
  - b) (Optional) In the Description field, enter a description of the DHCP relay policy, if necessary.  
The description can be up to 128 alphanumeric characters.
  - c) Expand **Providers**.  
The **Create DHCP Provider** dialog box appears.
  - d) In the **Create DHCP Provider** dialog box, in the **EPG Type** field, click the appropriate radio button depending upon where the DHCP server is connected.  
The options for the EPG type that you choose varies, depending on the EPG type.
    - If choose **Application EPG** as the EPG type, the following options appear in the **Application EPG** area:
      - In the **Tenant** field, from the drop-down list, choose the tenant. (*infra*)
      - In the **Application Profile** field, from the drop-down list, choose the application. (*access*)
      - In the **EPG** field, from the drop-down list, choose the EPG. (*default*)
    - If choose **L2 External Network** as the EPG type, the following options appear in the **L2 External Network** area:
      - In the **Tenant** field, from the drop-down list, choose the tenant.
      - In the **L2 Out** field, from the drop-down list, choose the L2 Out.
      - In the **External Network** field, from the drop-down list, choose the external network.

- If choose **L3 External Network** as the EPG type, the following options appear in the **L3 External Network** area:
    - In the **Tenant** field, from the drop-down list, choose the tenant.
    - In the **L3 Out** field, from the drop-down list, choose the L3 Out.
    - In the **External Network** field, from the drop-down list, choose the external network.
  - If choose **DN** as the EPG type, enter the distinguished name of the target endpoint group.
- e) In the **DHCP Server Address** field, enter the IP address for the DHCP server.
- f) In the **DHCP Server Preference** field, select the administrative preference value for this provider.

The **DHCP Server Preference** field is available beginning with release 5.2(4). Using the value in this field, the leaf switch determines whether to route the DHCP relay packets from the client VRF or the server VRF. For more information, see [About the DHCP Server Preference Field, on page 132](#).

- **None:** This is the default option, which reflects the behavior prior to release 5.2(4). By choosing the **None** option, the switch will always route the DHCP relay packet from the client VRF. If used for inter-VRF DHCP relay, a shared services contract is required to leak the server VRF network into the client VRF.
- **Use Server VRF:** This option reflects new behavior introduced in release 5.2(4). By choosing the **Use Server VRF** option, the switch routes the DHCP relay packets from the server VRF, regardless of whether there is a contract or there is no contract between the EPG where the DHCP client is and the EPG where the DHCP server is (or the Layer 3 external of the L3Out through which the DHCP server is reachable).

For inter-VRF configurations, when you choose the **Use Server VRF** option in the **DHCP Server Preference** field, the server subnet route is programmed in the server VRF on the client leaf switch for route lookup. The DHCP process on the client leaf switch then sends the DHCP relay packets via the server VRF. Because of this, the server VRF must be deployed with at least one IP address on all leaf switches where the client bridge domains are deployed.

- g) Click **OK**.
- You are returned to the **Create DHCP Relay Policy** window.
- h) Click **Submit**.
- The DHCP relay policy is created.

#### Step 5

**Step 6** In the **Navigation** pane, expand **Networking > Bridge Domains > default > DHCP Relay Labels**.

**Step 7** Right-click **DHCP Relay Labels**, and click **Create DHCP Relay Label**.

**Step 8** In the **Create DHCP Relay Label** dialog box, perform the following actions:

- a) In the **Scope** field, click the tenant radio button.  
This action displays, in the **Name** field drop-down list, the DHCP relay policy created earlier.
- b) In the **Name** field, from the drop-down list, choose the name of the DHCP policy created (DhcpRelayP) or create a new relay policy by choosing **Create DHCP Relay Policy**.
- c) In the **DHCP Option Policy**, select an existing option policy, or create a new one by choosing **Create DHCP Option Policy**.

For invoking option 79, select the DHCP option policy earlier created with 79 as the **ID**.



If you are creating a new option policy, in the **Create DHCP Option Policy** window, in the **Options** pane, ensure to enter the **ID** as 79.

d) Click **Submit**.

The DHCP server is associated with the bridge domain.

**Step 9** In the **Navigation** pane, expand **Networking > Bridge Domains > default > DHCP Relay Labels** to view the DHCP server created.

## Configuring Option 79 Using REST API

To configure Option 79 for a DHCP option policy using REST API:

POST URL: `https://apic-ip-address/api/mo/uni.xml`

```
<dhcpOptionPol dn="uni/tn-dhcp_client/dhcpoptpol-dhcp_option_policy" name="dhcp_option_policy"
 status="">
<dhcpOption data="" id="79" name="option_79"/>
</dhcpOptionPol>
```

## Configuring a DHCP Server Policy for the APIC Infrastructure Using the NX-OS Style CLI

- The port and the encapsulation used by the application Endpoint Group must belong to a physical or VM Manager (VMM) domain. If no such association with a domain is established, the APIC continues to deploy the EPG but raises a fault.
- Cisco APIC supports DHCP relay for both IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6. DHCPv6 relay will occur only if IPv6 is enabled on the fabric interface and one or more DHCPv6 relay servers are configured.

### Before you begin

Ensure that Layer 2 or Layer 3 connectivity is configured to reach the DHCP server address.

### Procedure

Configure DHCP server policy settings for the APIC infrastructure traffic.

#### Example:

#### DHCP Relay Policy for an Endpoint Group

```
apic1(config)# tenant infra
apic1(config-tenant)# template dhcp relay policy DhcpRelayP
apic1(config-tenant-template-dhcp-relay)# ip address 10.0.0.1 tenant infra application access epg default
apic1(config-tenant-template-dhcp-relay)# exit
apic1(config-tenant)# interface bridge-domain default
apic1(config-tenant-interface)# dhcp relay policy tenant DhcpRelayP
apic1(config-tenant-interface)# exit
```

#### Example:

### DHCP Relay Policy for Layer 3 Outside

```

ifav28-ifc2(config)# tenant dhcpTn
ifav28-ifc2(config-tenant)# template dhcp relay policy DhcpRelayPol
ifav28-ifc2(config-tenant-template-dhcp-relay)# ip address 11.1.1.11 tenant dhcpTn application ap epg serverEpg
ifav28-ifc2(config-tenant-template-dhcp-relay)# exit
ifav28-ifc2(config-tenant)# exit
ifav28-ifc2(config)# leaf 2001
ifav28-ifc2(config-leaf)# interface ethernet 1/4
ifav28-ifc2(config-leaf-if)# no switchport
ifav28-ifc2(config-leaf-if)# vrf member tenant dhcpTn vrf v1
ifav28-ifc2(config-leaf-if)# dhcp relay policy tenant DhcpRelayPol
ifav28-ifc2(config-leaf-if)# exit

```

## Configuring a DHCP Server Policy for the APIC Infrastructure Using the REST API

- This task is a prerequisite for users who want to create a vShield domain profile.
- The port and the encapsulation used by the application endpoint group must belong to a physical or VM Manager (VMM) domain. If no such association with a domain is established, the Cisco Application Policy Infrastructure Controller (APIC) continues to deploy the EPG but raises a fault.
- Cisco APIC supports DHCP relay for both IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6. DHCPv6 relay will occur only if IPv6 is enabled on the fabric interface and one or more DHCPv6 relay servers are configured.

### Before you begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

### Procedure

Configure the Cisco APIC as the DHCP server policy for the infrastructure tenant.

#### Note

This relay policy will be pushed to all the leaf ports that are connected hypervisors using the attach entity profile configuration. For details about configuring with attach entity profile, see the examples related to creating VMM domain profiles.

#### Example:

##### DHCP Relay Policy for EPG

```

<!-- api/policymgr/mo/.xml -->
<polUni>

POST https://apic-ip-address/api/mo/uni.xml

<fvTenant name="infra">

 <dhcpRelayP name="DhcpRelayP" owner="tenant">
 <dhcpRsProv tDn="uni/tn-infra/ap-access/epg-default" addr="10.0.0.1" />
 </dhcpRelayP>

```

```

 <fvBD name="default">
 <dhcpLbl name="DhcpRelayP" owner="tenant"/>
 </fvBD>

 </fvTenant>
</polUni>

```

**Example:****DHCP Relay Policy for Layer 3 Outside****Note**

You must specify DHCP Relay label under **l3extLIIfP** with an appropriate name and owner.

```

<polUni>
 <fvTenant name="dhcpTn">
 <l3extOut name="Out1" >
 <l3extLNodeP name="NodeP" >
 <l3extLIIfP name="Intf1">
 <dhcpLbl name="DhcpRelayPol" owner="tenant" />
 </l3extLIIfP>
 </l3extLNodeP>
 </l3extOut>
 </fvTenant>
</polUni>

```

POST <https://apic-ip-address/api/mo/uni.xml>

## Configuring a DNS Service Policy

A DNS policy is required to connect to external servers, for example AAA, RADIUS, vCenter, and services by hostname. A DNS service policy is a shared policy, so any tenant and VRF that uses this service must be configured with the specific DNS profile label. To configure a DNS policy for the ACI fabric, you must complete the following tasks:

- Ensure that the management EPG is configured for the DNS policy, otherwise this policy will not take into effect on the switches.




---

**Note** For the management EPG, only the default DNS policy is supported.

---

- Create a DNS profile (default) that contains the information about DNS providers and DNS domains.
- Associate the DNS profile (default or another DNS profile) name to a DNS label under the required tenant.

It is possible to configure a per-tenant, per-VRF DNS profile configuration. Additional DNS profiles can be created and applied to specific VRFs of specific tenants using the appropriate DNS label. For example, if you create a DNS profile with a name of acme, you can add a DNS label of acme to the appropriate **Networking > VRF** policy configuration in the tenants configuration.

## Configuring External Destinations with an In-Band DNS Service Policy

Configure the external destinations for the services as follows:

Source	In-Band Management	Out-of-Band Management	External Server Location
APIC	IP address or Fully Qualified domain name (FQDN)	IP address or FQDN	Anywhere
Leaf switches	IP address	IP address or FQDN <b>Note</b> The DNS policy must specify the out-of-band management EPG for reachability of the DNS server.	Anywhere
Spine switches	IP address	IP address or FQDN <b>Note</b> The DNS policy must specify the out-of-band management EPG for reachability of the DNS server.	Directly connected to a leaf switch

The following is a list of external servers:

- Call Home SMTP server
- Syslog server
- SNMP Trap destination
- Statistics Export destination
- Configuration Export destination
- Techsupport Export destination
- Core Export destination

The recommended guidelines are as follows:

- The external servers must be attached to the leaf access ports.
- Use in-band connectivity for the leaf switches to avoid extra cabling for the management port.
- Use out-of-band management connectivity for the spine switches. Connect this out-of-band network for spine switches to one of the leaf ports with in-band management virtual routing and forwarding (VRF) so that the spine switches and the leaf switches can reach the same set of external servers.
- Use IP addresses for the external servers.

## Dual Stack IPv4 and IPv6 DNS Servers

DNS servers have primary DNS records which can be A records (IPv4) or AAAA records (IPv6). Both A and AAAA records associate domain name with a specific IP address (IPv4 or IPv6).

The ACI fabric can be configured to use reputable public DNS servers that run on IPv4. These servers are able to resolve and respond with A record (IPv4) or AAAA record (IPv6).

In a pure IPv6 environment, the system administrators must use IPv6 DNS servers. The IPv6 DNS servers are enabled by adding them to `/etc/resolv.conf`.

A more common environment is to have dual-stack IPv4 and IPv6 DNS servers. In the dual-stack case, both IPv4 and IPv6 name servers are listed in `/etc/resolv.conf`. However, in a dual-stack environment, simply appending the IPv6 DNS servers to the list may cause a large delay in DNS resolutions. This is because the IPv6 protocol takes precedence by default, and it is unable to connect to the IPv4 DNS servers (if they are listed first in `/etc/resolv.conf`). The solution is to list IPv6 DNS servers ahead of IPv4 DNS servers. Also add “options single-request-reopen” to enable the same socket to be used for both IPv4 and IPv6 lookups.

Here is an example of `resolv.conf` in dual-stack IPv4 and IPv6 DNS servers where the IPv6 DNS servers are listed first. Also note the “single-request-reopen” option:

```
options single-request-reopen
nameserver 2001:4860:4680::8888
nameserver 2001:4860:4680::8844
nameserver 8.8.8.8
nameserver 8.8.4.4
```

## Dual-Stack IPv4 and IPv6 Environment

If the management network in the ACI fabric supports both IPv4 and IPv6, the Linux system application (glibc) will use the IPv6 network by default because `getaddrinfo()` will return IPv6 first.

Under certain conditions however, an IPv4 address may be preferred over an IPv6 address. The Linux IPv6 stack has a feature which allows an IPv4 address mapped as an IPv6 address using IPv6 mapped IPv4 address (`::ffff/96`). This allows an IPv6 capable application to use only a single socket to accept or connect both IPv4 and IPv6. This is controlled by the glibc IPv6 selection preference for `getaddrinfo()` in `/etc/gai.conf`.

In order to allow glibc to return multiple addresses when using `/etc/hosts`, “multi on” should be added to the `/etc/hosts` file. Otherwise, it may return only the first match.

If an application is not aware whether both IPv4 and IPv6 exist, it may not perform fallback attempts using different address families. Such applications may require a fallback implementation.

## Policy for Priority of IPv4 or IPv6 in a DNS Profile

The DNS profile supports version preference choices between IPv4 and IPv6. Using the user interface, you can enable your preference. IPv4 is the default.

The following is an example of a policy based configuration using Postman REST API:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/fabric/dnsp-default.xml -->
<dnsProfile dn="uni/fabric/dnsp-default" IPVerPreference="IPv6" childAction="" descr="" >
</dnsProfile>
```

The `gai.conf` settings control destination address selection. The file has a label table, precedence table, and an IPv4 scopes table. The changes for prioritizing IPv4 or IPv6 over the other need to go into the precedence

table entries. Given below are sample contents of the standard file as it is used in Linux systems for many flavors. A single line of precedence label in the file overrides any default settings.

The following is an example of a `gai.conf` to prioritize IPv4 over IPv6:

```
Generated by APIC
label ::1/128 0
label ::/0 1
label 2002::/16 2
label ::/96 3
label ::ffff:0:0/96 4
precedence ::1/128 50
precedence ::/0 40
precedence 2002::/16 30
precedence ::/96 20
For APICs preferring IPv4 connections, change the value to 100.
precedence ::ffff:0:0/96 10
```

## Configuring a DNS Service Policy to Connect with DNS Providers Using the GUI

### Before you begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

### Procedure

- 
- Step 1** On the menu bar, choose **FABRIC > Fabric Policies**. In the **Navigation** pane, expand **Global Policies > DNS Profiles**, and click the default DNS profile.
- Step 2** In the **Work** pane, in the **Management EPG** field, from the drop-down list, choose the appropriate management EPG (default (Out-of-Band)).
- Step 3** Expand **DNS Providers**, and perform the following actions:
- In the **Address** field, enter the provider address.
  - In the **Preferred** column, check the check box if you want to have this address as the preferred provider.  
You can have only one preferred provider.
  - Click **Update**.
  - (Optional) To add a secondary DNS provider, expand **DNS Providers**, and in the **Address** field, type the provider address. Click **Update**.
- Step 4** Expand **DNS Domains**, and perform the following actions:
- In the **Name** field, enter the domain name (cisco.com).
  - In the **Default** column, check the check box to make this domain the default domain.  
You can have only one domain name as the default.
  - Click **Update**.
  - (Optional) To add a secondary DNS domain, expand **DNS Domains**. In the **Address** field, enter the secondary domain name. Click **Update**.
- Step 5** Click **Submit**.  
The DNS server is configured.

- Step 6** On the menu bar, click **TENANTS > mgmt**.
- Step 7** In the **Navigation** pane, expand **Networking > VRF > oob**, and click **oob**.
- Step 8** In the **Work** pane, under **Properties**, in the **DNS labels** field, enter the appropriate DNS label (default). Click **Submit**.  
The DNS profile label is now configured on the tenant and VRF.
- 

## Configuring Custom Certificates

### Configuring Custom Certificate Guidelines

- Exporting a private key that is used to generate a Certificate Signing Request (CSR) on the Cisco Application Policy Infrastructure Controller (APIC) is not supported. If you want to use the same certificate on multiple servers through a wildcard in the Subject Alternative Name (SAN) field, such as "\*cisco.com," by sharing the private key that was used to generate the CSR for the certificate, generate the private key outside of Cisco Application Centric Infrastructure (ACI) fabric and import it to the Cisco ACI fabric.
- You must download and install the public intermediate and root CA certificates before generating a Certificate Signing Request (CSR). Although a root CA Certificate is not technically required to generate a CSR, Cisco requires the root CA certificate before generating the CSR to prevent mismatches between the intended CA authority and the actual one used to sign the CSR. The Cisco APIC verifies that the certificate submitted is signed by the configured CA.
- To use the same public and private keys for a renewed certificate generation, you must satisfy the following guidelines:
  - You must preserve the originating CSR as it contains the public key that pairs with the private key in the key ring.
  - The same CSR used for the originating certificate must be resubmitted for the renewed certificate if you want to re-use the public and private keys on the Cisco APIC.
  - Do not delete the original key ring when using the same public and private keys for the renewed certificate. Deleting the key ring will automatically delete the associated private key used with CSRs.
- Cisco ACI Multi-Site, VCPlugin, VRA, and SCVMM are not supported for certificate-based authentication.
- Only one SSL certificate is allowed per Cisco APIC cluster.
- You must disable certificate-based authentication before downgrading to release 4.0(1) from any later release.
- To terminate the certificate-based authentication session, you must log out and then remove the CAC card.
- The custom certificate configured for the Cisco APIC will be deployed to the leaf and spine switches. If the URL or DN that is used to connect to the fabric node is within the **Subject** or **Subject Alternative Name** field, the fabric node will be covered under the certificate.
- The Cisco APIC GUI can accept a certificate with a maximum size of 4k bytes.

- When a self-signed SSL certificate that you are using for HTTPS access expires, the certificate gets renewed automatically.

## Configuring a Custom Certificate for Cisco ACI HTTPS Access Using the GUI

**CAUTION: PERFORM THIS TASK ONLY DURING A MAINTENANCE WINDOW AS THERE IS A POTENTIAL FOR DOWNTIME.** The downtime affects access to the Cisco Application Policy Infrastructure Controller (APIC) cluster and switches from external users or systems and not the Cisco APIC to switch connectivity. The NGINX process on the switches will also be impacted, but that will be only for external connectivity and not for the fabric data plane. Access to the Cisco APIC, configuration, management, troubleshooting, and such will be impacted. The NGINX web server running on the Cisco APIC and switches will be restarted during this operation.

### Before you begin

Determine from which authority you will obtain the trusted certification so that you can create the appropriate Certificate Authority.

### Procedure

- 
- Step 1** On the menu bar, choose **Admin > AAA**.
- Step 2** In the **Navigation** pane, choose **Security**.
- Step 3** In the **Work** pane, choose **Public Key Management > Certificate Authorities > Create Certificate Authority**.
- Step 4** In the **Create Certificate Authority** dialog box, in the **Name** field, enter a name for the certificate authority.
- Step 5** In the **Certificate Chain** field, copy the intermediate and root certificates for the certificate authority that will sign the Certificate Signing Request (CSR) for the Cisco APIC.
- The certificate should be in Base64 encoded X.509 (CER) format. The intermediate certificate is placed before the root CA certificate. It should look similar to the following example:
- ```
-----BEGIN CERTIFICATE-----
<Intermediate Certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Root CA Certificate>
-----END CERTIFICATE-----
```
- Step 6** Click **Submit**.
- Step 7** In the **Navigation** pane, choose **Public Key Management > Key Rings**.
- Step 8** In the **Work** pane, choose **Actions > Create Key Ring**.
- The key ring enables you to manage a private key (imported from external device or internally generated on the Cisco APIC), a CSR generated by the private key, and the certificate signed via the CSR.
- Step 9** In the **Create Key Ring** dialog box, in the **Name** field, enter a name.
- Step 10** In the **Certificate** field, do not add any content if you will generate a CSR using the Cisco APIC through the key ring. Alternately, add the signed certificate content if you already have one that was signed by the CA from the previous steps by generating a private key and CSR outside of the Cisco APIC,
- Step 11** In the **Modulus** field, click the radio button for the desired key strength.

Step 12 In the **Certificate Authority** field, from the drop-down list, choose the certificate authority that you created earlier, then click **Submit**.

Step 13 In the **Private Key** field, do not add any content if you will generate a CSR using the Cisco APIC through the key ring. Alternately, add the private key used to generate the CSR for the signed certificate that you entered in step 10.

Note

Do not delete the key ring. Deleting the key ring will automatically delete the associated private key used with CSRs.

If you have not entered the signed certificate and the private key, in the **Work** pane, in the **Key Rings** area, the **Admin State** for the key ring created displays **Started**, waiting for you to generate a CSR. Proceed to step 14.

If you entered both the signed certificate and the private key, in the **Key Rings** area, the **Admin State** for the key ring created displays **Completed**. Proceed to step 23.

Step 14 In the **Navigation** pane, choose **Public Key Management > Key Rings > key_ring_name**.

Step 15 In the **Work** pane, choose **Actions > Create Certificate Request**.

Step 16 In the **Subject** field, enter the common name (CN) of the CSR.

You can enter the fully qualified domain name (FQDN) of the Cisco APICs using a wildcard, but in a modern certificate, we generally recommend that you enter an identifiable name of the certificate and enter the FQDN of all Cisco APICs in the **Alternate Subject Name** field (also known as the *SAN* – Subject Alternative Name) because many modern browsers expect the FQDN in the SAN field.

Step 17 In the **Alternate Subject Name** field, enter the FQDN of all Cisco APICs, such as "DNS:apic1.example.com,DNS:apic2.example.com,DNS:apic3.example.com" or "DNS:*example.com".

Alternatively, if you want SAN to match an IP address, enter the Cisco APICs' IP addresses with the following format:

```
IP:192.168.2.1
```

You can use DNS names, IPv4 addresses, or a mixture of both in this field. IPv6 addresses are not supported.

Step 18 Fill in the remaining fields as appropriate.

Note

Check the online help information available in the **Create Certificate Request** dialog box for a description of the available parameters.

Step 19 Click **Submit**.

Inside the same key ring, the **Associated Certificate Request** area is now displayed with the **Subject**, **Alternate Subject Name** and other fields you entered in the previous steps along with the new field **Request**, which contains the content of the CSR that is tied to this key ring. Copy the content from the **Request** field to submit the content to the same certificate authority that is tied to this key ring for signing.

Step 20 In the **Navigation** pane, choose **Public Key Management > Key Rings > key_ring_name**.

Step 21 In the **Work** pane, in the **Certificate** field, paste the signed certificate that you received from the certificate authority.

Step 22 Click **Submit**.

Note

If the CSR was not signed by the Certificate Authority indicated in the key ring, or if the certificate has MS-DOS line endings, an error message is displayed and the certificate is not accepted. Remove the MS-DOS line endings.

The key is verified, and in the **Work** pane, the **Admin State** changes to **Completed** and is now ready for use in the HTTP policy.

Step 23 On the menu bar, choose **Fabric > Fabric Policies**.

Step 24 In the Navigation pane, choose **Policies > Pod > Management Access > default**.

Step 25 In the **Work** pane, in the **Admin Key Ring** drop-down list, choose the desired key ring.

Step 26 (Optional) For Certificate based authentication, in the **Client Certificate TP** drop-down list, choose the previously created Local User policy and click **Enabled** for **Client Certificate Authentication state**.

Step 27 Click **Submit**.

All web servers restart. The certificate is activated, and the non-default key ring is associated with HTTPS access.

What to do next

You must remain aware of the expiration date of the certificate and take action before it expires. To preserve the same key pair for the renewed certificate, you must preserve the CSR as it contains the public key that pairs with the private key in the key ring. Before the certificate expires, the same CSR must be resubmitted. Do not delete or create a new key ring as deleting the key ring will delete the private key stored internally on the Cisco APIC.

Provisioning Fabric Wide System Settings

Configuring APIC In-Band or Out-of-Band Connectivity Preferences

This topic describes how to toggle between in-band and out-of-band connectivity on the APIC server for management access to devices such as authentication servers or SNMP servers external to the ACI fabric. Enabling **inband** executes in-band management connectivity between the APIC server to external devices through leaf switches on the ACI fabric. Enabling **ooband** executes out-of-band management connectivity between the APIC server to external devices through connections external to the ACI fabric.

Before you begin

Configure in-band and out-of-band management networks. For more information, see *Management* in the *Cisco APIC Basic Configuration Guide, Release 3.x*.

Procedure

Step 1 On the menu bar, click **System > System Settings**.

Step 2 On the Navigation bar, click **APIC Connectivity Preferences**.

Step 3 To enable the policy, click **inband** or **ooband**.

Step 4 Click **Submit**.

Configure Quota Management Policies

Starting in the Cisco Application Policy Infrastructure Controller (APIC) Release 2.3(1), there are limits on number of objects a tenant admin can configure. This enables the admin to limit the number of managed objects that can be added globally across tenants.

This feature is useful when you want to limit any tenant or group of tenants from exceeding ACI maximums per leaf or per fabric or unfairly consuming a majority of available resources, potentially affecting other tenants on the same fabric.

Procedure

- Step 1** On the menu bar, click **System > System Settings**.
 - Step 2** Right-click **Quota** and choose **Create Quota Configuration**.
 - Step 3** In the **Class** field, choose the object type to limit with the quota.
 - Step 4** In the **Container Dn** field, enter the distinguished name (DN) that describes the class.
 - Step 5** In the **Exceed Action** field, choose either **Fail Transaction Action** or **Raise Fault Action**.
 - Step 6** In the **Max Number** field, enter the maximum number of the managed objects that can be created after which the exceed action will be applied.
 - Step 7** Click **Submit**.
-

Create an Enforced BD Exception List

This topic describes how to create a global exception list of subnets which are not subject to an enforced bridge domain. With the Enforced BD feature configured, the endpoints in a subject endpoint group (EPG) can only ping subnet gateways within the associated bridge domain.

The exception IP addresses can ping all of the BD gateways across all of your VRFs.

A loopback interface configured for an L3Out does not enforce reachability to the IP address that is configured for the subject loopback interface.

When an eBGP peer IP address exists in a different subnet than the subnet of the L3Out interface, the peer subnet must be added to the allowed exception subnets. Otherwise, eBGP traffic is blocked because the source IP address exists in a different subnet than the L3Out interface subnet.

Before you begin

Create an enforced bridge domain (BD).

Procedure

-
- Step 1** On the menu bar, click **System > System Settings**.
 - Step 2** Click **BD Enforced Exception List**.
 - Step 3** Click the + on **Exception List**.
 - Step 4** Add the IP address and network mask for the subnet that can ping any subnet gateway.
 - Step 5** Repeat to add more subnets that are exceptions to the enforced bridge domain.
 - Step 6** Click **Submit**.
-

Create a BGP Route Reflector Policy and Route Reflector Node Endpoints

This topic describes how to create ACI fabric route reflectors, which use multiprotocol BGP (MP-BGP) to distribute external routes within the fabric. To enable route reflectors in the ACI fabric, the fabric administrator must select the spine switches that will be the route reflectors, and provide the autonomous system (AS) number. Once route reflectors are enabled in the ACI fabric, administrators can configure connectivity to external networks.

Before you begin

Required:

- To connect external routers to the ACI fabric, the fabric infrastructure administrator must configure spine nodes as Border Gateway Protocol (BGP) route reflectors.
- For redundancy purposes, more than one spine is configured as a router reflector node (one primary and one secondary reflector).

Procedure

-
- Step 1** To create a BGP Route Reflector policy, perform the following steps:
 - a) On the menu bar, click **System > System Settings**.
 - b) Click **BGP Route Reflector**.
 - c) Enter the Autonomous System Number.
 - d) Click the + on **Route Reflector Nodes**.
 - e) Enter the spine route reflector node ID endpoint, and click **Submit**.
 - Step 2** To create external route reflector node endpoints, perform the following steps:
 - a) Click the + on **External Route Reflector Nodes**.
 - b) Choose the spine to serve as external route reflector node endpoint.
 - c) If this is a site managed by Multi-Site, you can also specify an intersite spine route reflector.
 - d) Click **Submit**.
-

Configure a Fabric Wide Control Plane MTU Policy

This topic describes how to create a fabric-wide Control Plane (CP) MTU policy, that sets the global MTU size for control plane packets sent by the nodes (switches) in the fabric.

In a multipod topology, the MTU setting for the fabric external ports must be greater than or equal to the CP MTU value set. Otherwise, the fabric external ports might drop the CP MTU packets.



Note

- If you set the L3Out Interface Profile to inherit the MTU from the IPN, it will be 9150. If you want the MTU to be used across the IPN to be 9216, you must explicitly configure it in the L3Out Interface Profile (at **Tenants > *tenant-name* > Networking > External Routed Networks > Create Routed Outside > Nodes and Interface Protocol Profiles > Create Node Profile > Create Interface Profile**).
- Cisco APIC will always establish a TCP connection to fabric switches with an MTU of 1496 bytes (TCP MSS 1456) regardless of the CP-MTU setting. The IPN network for remote pods and remote leaf switches must support at least 1500 byte MTU for fabric discovery.

If you change the IPN or CP MTU, Cisco recommends changing the CP MTU value first, then changing the MTU value on the spine of the remote pod. This reduces the risk of losing connectivity between the pods due to MTU mismatch.

Procedure

- Step 1** On the menu bar, click **System > System Settings**.
- Step 2** Click **Control Plane MTU**.
- Step 3** Enter the MTU for fabric ports.
- Step 4** Click **Submit**.

Configure Endpoint Loop Protection

The endpoint loop protection policy specifies how loops detected by frequent MAC moves are handled. To configure EP loop protection perform the following steps:

Procedure

- Step 1** On the menu bar, click **System > System Settings**.
- Step 2** Click **Endpoint Controls**.
- Step 3** Click the **Ep Loop Protection** tab.
- Step 4** To enable the policy, click **Enabled** in the **Administrative State** field.
- Step 5** Optional. Set the loop detection interval, which specifies the time to detect a loop. The interval range is from 30 to 300 seconds. The default setting is 60 seconds.

Step 6 Set the loop detection multiplication factor, which is the number of times a single EP moves between ports within the loop detection interval. The range is from 1 to 255. The default is 4.

Step 7 Choose the action to take when detecting a loop.

The action can be:

- **BD Learn Disable**
- **Port Disable**

The default is **Port Disable**.

Step 8 Click **Submit**.

Rogue Endpoint Control Policy

About the Rogue Endpoint Control Policy

A rogue endpoint attacks leaf switches through frequently, repeatedly injecting packets on different leaf switch ports and changing 802.1Q tags (thus, emulating endpoint moves) causing learned class and EPG port changes. Misconfigurations can also cause frequent IP and MAC address changes (moves).

Such rapid movement in the fabric causes significant network instability, high CPU usage, and in rare instances, endpoint mapper (EPM) and EPM client (EPMC) crashes due to significant and prolonged messaging and transaction service (MTS) buffer consumption. Also, such frequent moves may result in the EPM and EPMC logs rolling over very quickly, hampering debugging for unrelated endpoints.

The rogue endpoint control feature addresses this vulnerability by quickly:

- Identifying such rapidly moving MAC and IP endpoints.
- Stopping the movement by temporarily making endpoints static, thus quarantining the endpoint.
- Prior to 3.2(6) release: Keeping the endpoint static for the **Rogue EP Detection Interval** and dropping the traffic to and from the rogue endpoint. After this time expires, deleting the unauthorized MAC or IP address.
- In the 3.2(6) release and later: Keeping the endpoint static for the **Rogue EP Detection Interval** (this feature no longer drops the traffic). After this time expires, deleting the unauthorized MAC or IP address.
- Generating a host tracking packet to enable the system to re-learn the impacted MAC or IP address.
- Raising a fault to enable corrective action.

The rogue endpoint control policy is configured globally and, unlike other loop prevention methods, functions at the level of individual endpoints (IP and MAC addresses). It does not distinguish between local or remote moves; any type of interface change is considered a move in determining if an endpoint should be quarantined.

The rogue endpoint control feature is disabled by default.

Limitations of the Rogue Endpoint Control Policy

The following limitations apply when using a rogue endpoint control policy:

- Changing rogue endpoint control policy parameters will not affect existing rogue endpoints.

- If a rogue endpoint is enabled, loop detection and bridge domain move frequency will not take effect.
- Disabling the rogue endpoint feature clears all rogue endpoints.
- The endpoint mapper (EPM) has value limits for rogue endpoint parameters. If you set the parameter values outside of this range, the Cisco APIC raises a fault for each mismatched parameter.
- Support for rogue endpoint detection is limited to endpoints that are connected to the fabric and not to those that are connected to remote leaf nodes.
- The rogue endpoint feature can be used within each site of a Cisco ACI Multi-Site deployment to help with misconfigurations of servers that cause an endpoint to move within the site. The rogue endpoint feature is not designed for scenarios where the endpoint may move between sites.
- You must disable rogue endpoint control before you upgrade to or from Cisco APIC release 4.1.

Configuring the Rogue Endpoint Control Policy Using the GUI

You can configure the **Rogue EP Control** policy for the fabric to detect and delete unauthorized endpoints using the Cisco Application Policy Infrastructure Controller (Cisco APIC) GUI. This topic also includes the steps to clear rogue endpoints on a leaf switch, ad-hoc.

Procedure

-
- Step 1** On the menu bar, click **System > System Settings**.
- Step 2** In the Navigation pane, choose **Endpoint Controls**
- Step 3** In the Work pane, choose the **Rogue EP Control** tab.
- Step 4** Set the **Administrative State** to **Enabled**.
- Step 5** Set the **Rogue EP Detection Interval**, **Rogue EP Detection Multiplication Factor**, and **Hold Interval (sec)** to the desired values.
- **Rogue EP Detection Interval:** Sets the rogue endpoint detection interval, which specifies the time to detect rogue endpoints. Valid values are from 0 to 65535 seconds. The default is 60.
 - **Rogue EP Detection Multiplication Factor:** Sets the rogue endpoint detection multiplication factor for determining if an endpoint is unauthorized. If the endpoint moves more times than this number, within the endpoint detection interval, the endpoint is declared rogue. Valid values are from 2 to 10. The default is 6.
 - **Hold Interval (sec):** Interval in seconds after the endpoint is declared rogue, where it is kept static so learning is prevented and the traffic to and from the rogue endpoint is dropped. After this interval, the endpoint is deleted. Prior to the 5.2(3) release, the valid values are from 1800 to 3600 seconds. Beginning with the 5.2(3) release, the valid values are from 300 to 3600 seconds. The default is 1800.
- Step 6** (Optional) To clear rogue endpoints on a leaf switch, perform the following steps:
- a) On the Cisco APIC menu bar, click **Fabric > Inventory**.
 - b) On the Navigation bar, expand the Pod and click the leaf switch where you want to clear rogue endpoints.
 - c) When the leaf switch summary appears in the work pane, right-click the leaf switch name in the Navigation bar, and choose **Clear Rogue Endpoints**.

d) Click **Yes**.

Configure the Rogue Endpoint Control Policy Using the NX-OS Style CLI

You can configure the rogue endpoint control policy for the fabric to detect and delete unauthorized endpoints using the NX-OS style CLI.

Procedure

Step 1 Enter global configuration mode.

Example:

```
apic1# configure
```

Step 2 Enable the global rogue endpoint control policy.

Example:

```
apic1(config)# endpoint rogue-detect enable
```

Step 3 Set the hold interval.

The hold interval is a period of time in seconds after the endpoint is declared rogue that the endpoint is kept static so that learning is prevented, and the traffic to and from the endpoint is dropped. After this interval, the endpoint is deleted. In the 5.2(2) release and earlier, the valid values are from 1800 to 3600 seconds. In the 5.2(3) release and later, the valid values are from 300 to 3600 seconds. The default is 1800.

Example:

```
apic1(config)# endpoint rogue-detect hold-interval 1800
```

Step 4 Set the detection interval.

The detection interval is a period of time in seconds during which rogue endpoint control counts the number of moves for an endpoint. If the count during this interval exceeds the value specified by the detection multiplication factor, the endpoint is declared rogue. Valid values are from 0 to 65535 seconds. The default is 60.

Example:

```
apic1(config)# endpoint rogue-detect interval 60
```

Step 5 Set the detection multiplication factor.

If an endpoint moves more times than the value specified by the detection multiplication factor during a period of time specified by the detection interval, the endpoint is declared rogue. Valid values are from 2 to 10. The default is 6.

Example:

```
apic1# endpoint rogue-detect factor 6
```

Configure the Rogue Endpoint Control Policy Using the REST API

You can configure the rogue endpoint control policy for the fabric to detect and delete unauthorized endpoints using the REST API.

Procedure

Step 1 To configure the rogue endpoint control policy, send a post with XML similar to the following example:

Example:

```
<polUni>
  <infraInfra>
    <epControlP name="default" adminSt="enabled" holdIntvl="1800"
    rogueEpDetectIntvl="60" rogueEpDetectMult="6"/>
  </infraInfra>
</polUni>
```

- **adminSt:** The administrative state of rogue endpoint control. Specify *enable* to enable rogue endpoint control.
- **holdIntvl:** Rogue endpoint hold interval. The hold interval is a period of time in seconds after the endpoint is declared rogue that the endpoint is kept static so that learning is prevented, and the traffic to and from the endpoint is dropped. After this interval, the endpoint is deleted. In the 5.2(2) release and earlier, the valid values are from 1800 to 3600 seconds. In the 5.2(3) release and later, the valid values are from 300 to 3600 seconds. The default is 1800.
- **rogueEpDetectIntvl:** Rogue endpoint detection interval. The detection interval is a period of time in seconds during which rogue endpoint control counts the number of moves for an endpoint. If the count during this interval exceeds the value specified by the detection multiplication factor, the endpoint is declared rogue. Valid values are from 0 to 65535 seconds. The default is 60.
- **rogueEpDetectMult:** Rogue endpoint detection multiplication factor. If an endpoint moves more times than the value specified by the detection multiplication factor during a period of time specified by the detection interval, the endpoint is declared rogue. Valid values are from 2 to 10. The default is 6.

Step 2 You can revert this feature's behavior so that it once again drops the traffic to and from rogue endpoints by sending a post with XML similar to the following example:

Example:

```
<infraImplicitSetPol rogueModeAction="quarantine-fault-and-drop" infraDn="uni/infra"/>
```

About the Rogue/COOP Exception List

The rogue/COOP exception list enables you to specify the MAC address of endpoints for which you want to have a higher tolerance for endpoint movement with rogue endpoint control before the endpoints get marked as rogue. Endpoints in the rogue/COOP exception list get marked as rogue only if they move 3,000 or more times within 10 minutes. After an endpoint is marked as rogue, the endpoint is kept static to prevent learning. The rogue endpoint is deleted after 30 seconds.

Beginning with the Cisco Application Policy Infrastructure Controller (APIC) 6.0(3) release, you can create a global rogue/COOP exception list, which excludes a MAC address from rogue endpoint control on all the

bridge domains on which the MAC address is discovered, and you can create a rogue/COOP exception list for L3Outs. You can also exclude all MAC addresses for a bridge domain or L3Out. This simplifies creating the exception list when you want to make an exception for every MAC address; you do not need to enter each address individually.

Guidelines and Limitations for the Rogue/COOP Exception List

The following guidelines and limitations apply when using the rogue/COOP exception list:

- The MAC address exception list feature works on Layer 2 bridge domains (bridge domains without IP routing enabled). This is because on a Layer 3 bridge domain (a bridge domain with IP routing enabled), if there are IP addresses moving along with MAC addresses, the IP addresses would end up being marked rogue first and both the IP and MAC address would then be quarantined.
- For a Layer 3 bridge domain, disable dataplane IP address learning per subnet for specific IP addresses that you want to exclude from rogue endpoint control.

For information about the dataplane IP address learning per subnet feature, see the *Cisco APIC Layer 3 Networking Configuration Guide*.

- You are expected to have a full understanding of what kind of MAC addresses are being added to this list. You are responsible for ensuring that the MAC addresses in this list do not contribute to excessive moves in the overall fabric or per leaf switch.
- You can add up to 100 MAC addresses to the per-bridge domain exception list, fabric-wide. Beginning in the 6.0(3) release you can also exclude all MAC addresses of a given bridge domain from rogue endpoint control.
- Beginning in the 6.0(3) release, you can add 100 external L3Out MAC addresses to the exception list, fabric-wide. You can also exclude all MAC addresses of a L3Out SVI bridge domain from rogue endpoint control.
- Beginning in the 6.0(3) release, you can add up to 6,000 MAC addresses to the global exception list as long as the total count of endpoints for those MAC addresses is within 6,000. The same MAC address discovered on a number of bridge domains counts as that same number of bridge domains against the total of 6,000 endpoints. For example, if a configured MAC address appears as a discovered MAC endpoint on 10 bridge domains, the MAC address counts as 10 endpoints.
- The exception list exemption on leaf switches applies only when rogue endpoint control is enabled. When rogue endpoint control is disabled, the MAC address exception list is used only by COOP dampening.
- The rogue/COOP exception list can only contain MAC addresses in a bridge domain and not IP addresses in a VRF instance. However, an IP address-only move can still cause the IP address to be marked as rogue if it meets the regular rogue endpoint control criteria.
- To mask out IP address rogue detection and marking based on data path traffic, you can use bridge domain subnet learn disable. Bridge domain subnet learn disable stops Cisco ACI from learning the IP address location at each move.

Configuring the Rogue/COOP Exception List While Creating a Bridge Domain Using the GUI

The following procedure configures the rogue/COOP exception list while you create a bridge domain:

Before you begin

- You must have a tenant for which you will create a bridge domain.
- Rogue endpoint control must be enabled. For the procedure about enabling rogue endpoint control, see [Configuring the Rogue Endpoint Control Policy Using the GUI, on page 151](#).

Procedure

-
- Step 1** In the desired tenant, create a bridge domain. In the menu bar, choose **Tenants** > *tenant_name*.
- Step 2** In the Navigation pane, choose **Networking** > **Bridge Domains**.
- Step 3** Right-click **Bridge Domains** and choose **Create Bridge Domain**.
- Step 4** In the **Create Bridge Domain** dialog, fill out the fields as desired for **STEP 1** > **MAIN** and **STEP 2** > **L3 Configurations**.
- Step 5** For **STEP 3** > **Advanced/Troubleshooting**, for the **Rogue/Coop Exception List**, click the +, enter the MAC address of an endpoint that you want to add to the list, and click **Update**.
- The format of the MAC address is AA:BB:CC:DD:EE:FF.
- a) Repeat this step for each endpoint that you want to add to the list.
- Step 6** (Optional): Beginning with the Cisco Application Policy Infrastructure Controller (APIC) 6.0(3) release, put a check in the **Enable Rogue Exception MAC for BD** check box to cause rogue endpoint control to ignore all MAC addresses from the bridge domain.
- Step 7** Fill out the remaining fields for **STEP 3** > **Advanced/Troubleshooting** as desired.
- Step 8** Click **Finish**.
-

Configuring the Rogue/COOP Exception List of an Existing Bridge Domain Using the GUI

The following procedure configures the rogue/COOP exception list of an existing bridge domain.

Before you begin

- You must have a tenant with a bridge domain.
- Rogue endpoint control must be enabled. For the procedure about enabling rogue endpoint control, see [Configuring the Rogue Endpoint Control Policy Using the GUI, on page 151](#).

Procedure

-
- Step 1** In the menu bar, choose **Tenants** > *tenant_name*.
- Step 2** In the Navigation pane, choose **Networking** > **Bridge Domains** > *bridge_domain_name*.
- Step 3** In the Work pane, choose **Policy** > **Advanced/Troubleshooting**.
- Step 4** For the **Rogue/Coop Exception List**, click the +, enter the MAC address of an endpoint that you want to add to the list, and click **Update**.

The format of the MAC address is AA:BB:CC:DD:EE:FF.

a) Repeat this step for each endpoint that you want to add to the list.

Step 5 (Optional): Beginning with the Cisco Application Policy Infrastructure Controller (APIC) 6.0(3) release, put a check in the **Enable Rogue Exception MAC for BD** check box to cause rogue endpoint control to ignore all MAC addresses from the bridge domain.

Step 6 Click **Submit**.

Configuring the Rogue/COOP Exception List of an Existing Bridge Domain Using the REST API

The following example REST API post adds a MAC address to the rogue/COOP exception list of an existing bridge domain:

```
https://apic1.myDomain.com/api/node/mo/uni/tn-tenant1.xml
<fvBD name="bd1">
  <fvRogueExceptionMac annotation="" descr="" mac="00:16:04:00:00:1"/>
</fvBD>
```

The following example REST API post adds a wildcard to the rogue/COOP exception list of an existing bridge domain, which creates an exception for all MAC addresses:

```
https://apic1.myDomain.com/api/node/mo/uni/tn-tenant1.xml
<fvBD name="bd1">
  <fvRogueExceptionMac annotation="" descr="" mac="FF:FF:FF:FF:FF:FF"/>
</fvBD>
```

About Max IP Address Flow Control

The 3.2(6) release adds the max IP address flow control feature, which identifies endpoints that are misbehaving and flags them as rogue based on the number of learned IP addresses that are associated with a MAC address. The Cisco Application Centric Infrastructure (ACI) fabric supports a maximum of 4,096 IP addresses on a MAC address. If a leaf switch learns more than 4,096 IP addresses that are associated with a MAC address, then the MAC address and all of the IP addresses are classified as rogue.

After the max IP address flow control feature identifies an endpoint as rogue, the endpoint is quarantined, a fault is raised in the APIC, and there will not be any further learning of new IP Addresses on this endpoint. The quarantine period is 1 hour. If the standard rogue feature is enabled, then the quarantine period is the same as the period configured by the standard rogue configuration.

While the rogue endpoint control policy feature (rogue due to moves) can be configured to be enabled or disabled, the max IP address flow control feature does not require explicit configuration to be enabled.

Prior to this feature, the ACI fabric identified an endpoint as rogue if it kept moving its location for a configurable number of times within a configurable time period. With this feature, the ACI fabric can identify an endpoint as rogue based on the number of moves or if a leaf switch learns more than 4,096 IP addresses on a MAC address.

Configure COOP

About COOP

Council of Oracle Protocol (COOP) is used to communicate the mapping information (location and identity) to the spine switch proxy. A leaf switch (the "citizen") forwards endpoint address information to the spine

switch (the "oracle") using Zero Message Queue (ZMQ). COOP running on the spine nodes will ensure all spine nodes maintain a consistent copy of endpoint address and location information and additionally maintain the distributed hash table (DHT) repository of endpoint identity to location mapping database.

COOP Endpoint Dampening

When malicious or erroneous behavior causes unnecessary endpoint updates, the COOP process can become overwhelmed, preventing the processing of valid endpoint updates. The rogue endpoint detection feature of the leaf switch can prevent many erroneous updates from reaching the spine switch. In cases where the rogue endpoint detection is inadequate, the COOP process invokes endpoint dampening. To relieve pressure on COOP, the spine switch asks all leaf switches to ignore updates from the misbehaving endpoint for a specified period. When this occurs, the dampening state of the endpoint is 'Freeze,' and a fault is generated.



Note COOP endpoint dampening is introduced and enabled by default in Cisco Application Policy Infrastructure Controller (APIC) release 4.2(3).

Detection criteria is based on the penalty value calculations that are based on the types of endpoint-related events, as shown in the following table:

Event	Penalty Value	Notes
Learn a new IP address	0	The new IP address is learned.
Learn an additional IP address	2	The additional IP address is learned with an existing endpoint MAC address.
Delete an IP address	50	The remote endpoint IP address gets deleted after the IP address is learned.
Learn the deleted IP address	50	Learn the remote endpoint IP address after the IP address is deleted.
Delete an IP address	400	Delete the local endpoint IP address after the IP address is learned.
Learn the deleted IP address	400	Learn the local endpoint IP address after the IP address is deleted.
Endpoint move	200	An endpoint moves to a different interface.
IP address move	200	An IP address moves to a different MAC address. The penalty is high for this event because it causes two route updates to BGP.
URIB programming	50	Spine switch tunnel interface status change (up/down) for an endpoint.

The penalty value is calculated per IP address and is decreased by 50% every five minutes. For example, if the penalty value of the endpoint is 4000 and number of IP addresses in the endpoint is 2, the penalty value per IP address is $4000/2 = 2000$. When the penalty value per IP address exceeds the critical threshold (4000),

the endpoint state is changed to **Critical** from **Normal**. If an endpoint stays in the **Critical** state more than five minutes or the penalty value per IP address exceeds the freeze threshold (10000), the endpoint state will become **Freeze** (dampening) and the update of the endpoint is ignored. When the penalty value per IP address drops below the reuse threshold (2500), the endpoint state becomes **Normal** (non-dampening). 10 minutes must elapse for the penalty value to be decreased by 75% ($10000 * 0.5 * 0.5 = 2500$). The thresholds are not user-configurable.

COOP Authentication

COOP data path communication provides high priority to transport using secured connections. To protect COOP messages from malicious traffic injection, the Cisco APIC and switches support COOP protocol authentication.

The COOP protocol supports two ZMQ authentication modes:

- **Strict mode:** COOP allows MD5 authenticated ZMQ connections only.
- **Compatible mode:** COOP accepts both MD5 authenticated and non-authenticated ZMQ connections for message transportation.

For additional information about COOP authentication, see the *Cisco APIC Security Configuration Guide*.

Viewing COOP Dampened Endpoints Using the GUI

Use this Cisco Application Policy Infrastructure Controller (APIC) GUI procedure to view all dampened endpoints in a spine node.

Procedure

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

Step 2 In the Navigation pane, expand the pod and the spine node.

Step 3 Expand **Protocols > COOP** and the **COOP** instance.

Step 4 Click **Endpoint Database** to display the endpoints.

Inspect the **Dampened State** column to find the dampened endpoints. The possible states are:

- **Normal:** The endpoint updates are normal.
 - **Critical:** Enough updates have been received that the endpoint could be moved into the freeze state. If an endpoint remains in the **Critical** state more than five minutes, the state changes to **Freeze**.
 - **Freeze:** Updates from this endpoint are currently being ignored due to frequent unnecessary updates. A fault has been generated.
-

Viewing COOP Dampened Endpoints Using the Switch CLI

Use this switch CLI procedure to view all dampened endpoints in a spine or leaf node.

Log in to the spine or leaf switch CLI and enter the following command:

```
show coop internal info repo ep dampening
```

Clearing COOP Dampened Endpoints Using the GUI

Use this Cisco Application Policy Infrastructure Controller (APIC) GUI procedure to clear and recover all dampened endpoints in a spine or leaf node. This operation must be executed on all spine switches and on the source leaf switch of the endpoint. If the dampened endpoint is still in the endpoint table on the leaf switch, the endpoint is published to the spine switch COOP database. If not, the dampened endpoint is deleted from the spine switch COOP database after two minutes.

Procedure

-
- Step 1** On the menu bar, click **Fabric > Inventory**.
 - Step 2** In the Navigation pane, expand the pod and the spine or leaf node.
 - Step 3** Right-click the node and choose **Clear Dampened Endpoints**.
 - Step 4** Click **Yes** to confirm the action.
-

Clearing a COOP Dampened Endpoint Using the Switch CLI

Use this procedure to clear and recover a dampened endpoint in a spine or leaf node. The procedure recovers a single endpoint whose dampening state is **Freeze**. This operation must be executed on all spine switches and on the source leaf switch of the endpoint.

Log in to the spine or leaf switch CLI and enter the following command:

```
clear coop internal info repo ep dampening key <bd> <mac>
```

Disabling COOP Endpoint Dampening Using the REST API

This procedure shows how to disable or enable COOP EP dampening using the APIC REST API.

COOP endpoint dampening is enabled by default, but in some situations it can be necessary to disable it. An example is when you are expecting many IP updates for a single MAC address and ignoring those updates could be disruptive to the network.

Use the following API, setting `disableEpDampening="true"` to disable COOP endpoint dampening.

```
<!-- api/policymgr/mo/.xml -->
<polUni>
  <infraInfra>
    <infraSetPol disableEpDampening="true"></infraSetPol>
  </infraInfra>
</polUni>
```

All nodes in the fabric will disable COOP endpoint dampening and will recover any existing endpoints whose dampened state is 'Freeze.'

Configuring COOP Authentication Using the APIC GUI

Procedure

- Step 1** On the menu bar, choose **System > System Settings**.
 - Step 2** In the **Navigation** pane, click on **COOP Group**.
 - Step 3** In the **Work** pane, under the **Policy Property** area in the **Type** field, choose the desired type from the **Compatible Type** and **Strict Type** options.
 - Step 4** Click **Submit**.
This completes the COOP authentication policy configuration.
-

Configuring COOP Authentication Using the Cisco NX-OS-Style CLI

Procedure

Configure the COOP authentication policy using the strict mode option.

Example:

```
apic1# configure
apic1(config)# coop-fabric
apic1(config-coop-fabric)# authentication type ?
compatible  Compatible type
strict      Strict type
apic101-apic1(config-coop-fabric)# authentication type strict
```

Configuring COOP Authentication Using the REST API

Procedure

Configure a COOP authentication policy.

In the example, the strict mode is chosen.

Example:

```
https://172.23.53.xx/api/node/mo/uni/fabric/pol-default.xml
```

```
<coopPol type="strict">
</coopPol>
```

Endpoint Listen Policy

About the Endpoint Listen Policy

You can configure an endpoint listen policy to detect untagged traffic that gets sent from anonymous endpoints to Cisco Application Centric Infrastructure (ACI) leaf switches that do not have an enforced policy. By default, if a policy is not deployed for a port, all endpoint traffic gets dropped on that port. If you configure an endpoint listen policy, this policy gets deployed on all leaf switch ports that do not have an existing enforced policy. The endpoint listen policy enables Cisco ACI to detect untagged traffic that arrives on those ports, such that Cisco ACI will know the MAC address or IP address of the anonymous endpoints. This allows the Cisco ACI administrator to decide in which EPG to put those endpoints. The Cisco Application Policy Infrastructure Controller (APIC) GUI displays all detected anonymous endpoints on the **Global Endpoints** configuration screen.



Note The endpoint listen policy is a beta feature. There is no guarantee that this feature will work as intended. Use at your own risk.

Configuring the Endpoint Listen Policy Using the GUI

This procedure configures an endpoint listen policy, which detects untagged traffic that gets sent from anonymous endpoints to Cisco Application Centric Infrastructure (ACI) leaf switches that do not have an enforced policy.



Note The endpoint listen policy is a beta feature. There is no guarantee that this feature will work as intended. Use at your own risk.

Procedure

-
- Step 1** On the menu bar, choose **System > System Settings**.
 - Step 2** In the Navigation pane, choose **Global Endpoints**.
 - Step 3** In the Work pane, put a check in the **End Point Listen Policy** check box.
 - Step 4** In the **End Point Listen Encap** drop-down list, choose **VLAN**.
 - Step 5** In the **End Point Listen Encap** text field, enter the VLAN ID. Valid values are from 1 to 4094. This should be a reserved VLAN encap, such that it can not be used by any EPGs.
 - Step 6** Click **Submit**.
-

Configure IP Aging

This topic describes how to enable an IP Aging policy. When enabled, the IP aging policy ages unused IPs on an endpoint.

When the Administrative State is enabled, the IP aging policy sends ARP requests (for IPv4) and neighbor solicitations (for IPv6) to track IPs on endpoints. If no response is given, the policy ages the unused IPs.

Procedure

-
- Step 1** On the menu bar, click **System > System Settings**.
 - Step 2** Click **Endpoint Controls**.
 - Step 3** Click the **Ip Aging** tab.
 - Step 4** To enable the policy, click **Enabled** in the **Administrative State** field.
-

What to do next

Create an End Point Retention policy, which is required, to specify the timer used for tracking IPs on endpoints. Navigate to **Tenants > tenant-name > > Policies > Protocol > End Point Retention**.

Disable Remote Endpoint Learning

This topic describes how to enable or disable IP end point learning.

The scope of this policy is fabric-wide. After configuration, a policy is pushed to each leaf switch as it comes up.

You should enable this policy in fabrics which include the Cisco Nexus 9000 series switches, 93128 TX, 9396 PX, or 9396 TX switches with the N9K-M12PQ uplink module, after all the nodes have been successfully upgraded to APIC Release 2.2(2x) or higher.

After any of the following configuration changes, you may need to manually flush previously learned IP endpoints:

- Remote IP endpoint learning is disabled
- The VRF is configured for ingress policy enforcement
- At least one Layer 3 interface exists in the VRF

To manually flush previously learned IP endpoints, enter the following command on both VPC peers: `vsh -c "clear system internal epm endpoint vrf <vrf-name> remote"`

To enable or disable IP end point learning, perform the following steps:

Procedure

-
- Step 1** On the menu bar, click **System > System Settings**.
 - Step 2** Click **Fabric Wide Setting**.
 - Step 3** Click the check box on **Disable Remote EP Learn**.
 - Step 4** Click **Submit**.
-

Globally Enforce Subnet Checks

This topic describes how to enable or disable subnet checking. When enabled, IP address learning is disabled outside of subnets configured in a VRF, for all other VRFs.

The scope of this policy is fabric-wide. After configuration, a policy is pushed to each leaf switch as it comes up.

Procedure

-
- Step 1** On the menu bar, click **System > System Settings**.
 - Step 2** Click **Fabric Wide Setting**.
 - Step 3** Click the check box on **Enforce Subnet Check**.
 - Step 4** Click **Submit**.
-

Reallocate a GIPo

This topic describes how to enable reallocating GIPos on non-stretched bridge domains to make room for stretched bridge domains.

With the introduction of Cisco ACI Multi-Site, there was a need to change the GIPo allocation scheme to provide the following benefits:

1. Minimize the number of bridge domains that have the same GIPo.
2. GIPos that are assigned to Cisco ACI Multi-Site stretched bridge domains do not overlap with GIPos that are assigned to non-stretched bridge domains.

To achieve this allocation, Cisco ACI introduced different pools whose sizes change based on amount of stretched and non-stretched bridge domains.

For a fresh installation of Cisco ACI, the Cisco APIC guarantees that both #1 and #2 are accomplished. During a Cisco ACI upgrade from a release prior to 2.3(1), the old schema is maintained to avoid fabric disruption because the existing GIPos might already be used for non-stretched bridge domains. As a result, Cisco ACI cannot guarantee that #2 is accomplished.

Enabling the **Reallocate GIPo** knob in Cisco APIC's fabric-wide setting policy causes Cisco APIC to re-allocate GIPos and use the newer allocation scheme. Enabling the knob is a one-time operation. Afterward, the GIPos will not overlap. This knob is relevant only in a Cisco ACI Multi-Site Orchestrator deployment if you upgrade from a release that is earlier than 2.3(1) to the 3.0(1) release or later.

The scope of this policy is fabric-wide. After configuration, a policy is pushed to each leaf switch as it comes up.

Procedure

-
- Step 1** On the menu bar, click **System > System Settings**.
 - Step 2** Click **Fabric Wide Setting**.

Step 3 Click the check box on **Reallocate Gipo**.

Step 4 Click **Submit**.

Globally Enforce Domain Validation

This topic describes how to enforce domain validation. When enabled, a validation check is performed when a static path is added, to determine if no domain is associated with an EPG.

The scope of this policy is fabric-wide. After configuration, a policy is pushed to each leaf switch as it comes up.

Procedure

Step 1 On the menu bar, click **System > System Settings**.

Step 2 Click **Fabric Wide Setting**.

Step 3 Click the check box on **Enforce Domain Validation**.

Step 4 Click **Submit**.

Enable OpFlex Client Authentication

This topic describes how to enable OpFlex client authentication for GOLF and Linux.

To deploy GOLF or Linux Opflex clients in an environment where the identity of the client cannot be guaranteed by the network, you can dynamically validate the client's identity based on a client certificate.



Note When you enable certificate enforcement, connectivity with any GOLF or Linux Opflex client that does not support client authentication is disabled.

The scope of this policy is fabric-wide. After configuration, a policy is pushed to each leaf switch as it comes up.

Procedure

Step 1 On the menu bar, click **System > System Settings**.

Step 2 Click **Fabric Wide Setting**.

Step 3 Click the check box on **OpFlex Client Authentication** to enable or disable enforcing client certificate authentication for GOLF and Linux Opflex clients.

Step 4 Click **Submit**.

Fabric Load Balancing

The Cisco Application Centric Infrastructure (ACI) fabric provides several load balancing options for balancing the traffic among the available uplink links. This topic describes load balancing for leaf to spine switch traffic.

Static hash load balancing is the traditional load balancing mechanism used in networks where each flow is allocated to an uplink based on a hash of its 5-tuple. This load balancing gives a distribution of flows across the available links that is roughly even. Usually, with a large number of flows, the even distribution of flows results in an even distribution of bandwidth as well. However, if a few flows are much larger than the rest, static load balancing might give suboptimal results.

Cisco ACI fabric Dynamic Load Balancing (DLB) adjusts the traffic allocations according to congestion levels. It measures the congestion across the available paths and places the flows on the least congested paths, which results in an optimal or near optimal placement of the data.

DLB can be configured to place traffic on the available uplinks using the granularity of flows or flowlets. Flowlets are bursts of packets from a flow that are separated by suitably large gaps in time. If the idle interval between two bursts of packets is larger than the maximum difference in latency among available paths, the second burst (or flowlet) can be sent along a different path than the first without reordering packets. This idle interval is measured with a timer called the flowlet timer. Flowlets provide a higher granular alternative to flows for load balancing without causing packet reordering.

DLB modes of operation are aggressive or conservative. These modes pertain to the timeout value used for the flowlet timer. The aggressive mode flowlet timeout is a relatively small value. This very fine-grained load balancing is optimal for the distribution of traffic, but some packet reordering might occur. However, the overall benefit to application performance is equal to or better than the conservative mode. The conservative mode flowlet timeout is a larger value that guarantees packets are not to be re-ordered. The tradeoff is less granular load balancing because new flowlet opportunities are less frequent. While DLB is not always able to provide the most optimal load balancing, it is never worse than static hash load balancing.



Note Although all Nexus 9000 Series switches have hardware support for DLB, the DLB feature is not enabled in the current software releases for second generation platforms (switches with EX, FX, and FX2 suffixes).

The Cisco ACI fabric adjusts traffic when the number of available links changes due to a link going off-line or coming on-line. The fabric redistributes the traffic across the new set of links.

In all modes of load balancing, static or dynamic, the traffic is sent only on those uplinks or paths that meet the criteria for equal cost multipath (ECMP); these paths are equal and the lowest cost from a routing perspective.

Dynamic Packet Prioritization (DPP), while not a load balancing technology, uses some of the same mechanisms as DLB in the switch. DPP configuration is exclusive of DLB. DPP prioritizes short flows higher than long flows; a short flow is less than approximately 15 packets. Because short flows are more sensitive to latency than long ones, DPP can improve overall application performance.

In the 6.0(1) and 6.0(2) releases, for intra-leaf switch traffic, all DPP-prioritized traffic is marked CoS 0 regardless of a custom QoS configuration. For inter-leaf switch traffic, all DPP-prioritized traffic is marked CoS 3 regardless of a custom QoS configuration. Beginning with the 6.0(3) release, for intra-leaf switch and inter-leaf switch traffic, all DPP-prioritized traffic is marked CoS 0 regardless of a custom QoS configuration.

GPRS tunneling protocol (GTP) is used mainly to deliver data on wireless networks. Cisco Nexus switches are placed in Telecom Datacenters. When packets are being sent through Cisco Nexus 9000 switches in a datacenter, traffic needs to be load-balanced based on the GTP header. When the fabric is connected with an

external router through link bundling, the traffic is required to be distributed evenly between all bundle members (For example, Layer 2 port channel, Layer 3 ECMP links, Layer 3 port channel, and L3Out on the port channel). GTP traffic load balancing is performed within the fabric as well.

To achieve GTP load balancing, Cisco Nexus 9000 Series switches use 5-tuple load balancing mechanism. The load balancing mechanism takes into account the source IP, destination IP, protocol, Layer 4 resource and destination port (if traffic is TCP or UDP) fields from the packet. In the case of GTP traffic, a limited number of unique values for these fields restrict the equal distribution of traffic load on the tunnel.

To avoid polarization for GTP traffic in load balancing, a tunnel endpoint identifier (TEID) in the GTP header is used instead of a UDP port number. Because the TEID is unique per tunnel, traffic can be evenly load balanced across multiple links in the bundle.

The GTP load balancing feature overrides the source and destination port information with the 32-bit TEID value that is present in GTPU packets.

GTP tunnel load balancing feature adds support for:

- GTP with IPv4/IPv6 transport header on physical interface
- GTPU with UDP port 2152

The Cisco ACI fabric default configuration uses a traditional static hash. A static hashing function distributes the traffic between uplinks from the leaf switch to the spine switch. When a link goes down or comes up, traffic on all links is redistributed based on the new number of uplinks.

Leaf/Spine Switch Dynamic Load Balancing Algorithms

The following table provides the default non-configurable algorithms used in leaf/spine switch dynamic load balancing:

Table 6: Cisco ACI Leaf/Spine Switch Dynamic Load Balancing

Traffic Type	Hashing Data Points
Leaf/Spine IP unicast	<ul style="list-style-type: none"> • Source MAC address • Destination MAC address • Source IP address • Destination IP address • Protocol type • Source Layer 4 port • Destination Layer 4 port • Segment ID (VXLAN VNID) or VLAN ID
Leaf/Spine Layer 2	<ul style="list-style-type: none"> • Source MAC address • Destination MAC address • Segment ID (VXLAN VNID) or VLAN ID

Creating a Load Balancer Policy Using the Cisco APIC GUI

This topic describes how to configure the default Load Balancer policy.

The load balancing policy options balance traffic among the available uplink ports. Static hash load balancing is the traditional load balancing mechanism used in networks where each flow is allocated to an uplink based on a hash of its 5-tuple. This load balancing gives a distribution of flows across the available links that is roughly even. Usually, with a large number of flows, the even distribution of flows results in an even distribution of bandwidth as well. However, if a few flows are much larger than the rest, static load balancing might give suboptimal results.

Procedure

Step 1 On the menu bar, click **System** > **System Settings**.

Step 2 Click **Load Balancer**.

Step 3 Choose the **Dynamic Load Balancing Mode**.

The dynamic load balancer (DLB) mode adjusts the traffic allocations according to congestion levels. It measures the congestion across the available paths and places the flows on the least congested paths, which results in an optimal or near optimal placement of the data. DLB can be configured to place traffic on the available uplinks using the granularity of flows or of flowlets. Flowlets are bursts of packets from a flow that are separated by intervals. The mode can be **Aggressive**, **Conservative**, or **Off** (the default).

Step 4 Enable or disable **Dynamic Packet Prioritization** by choosing **On** or **Off** (the default).

Dynamic Packet Prioritization (DPP) prioritizes short flows higher than long flows; a short flow is less than approximately 15 packets. Short flows are more sensitive to latency than long ones. DPP can improve overall application performance.

Step 5 Choose the Load Balancing Mode. The mode can be **Link Failure** or **Traditional** (the default).

The load balancer administrative state. In all modes of load balancing, static or dynamic, the traffic is sent only on those uplinks or paths that meet the criteria for equal cost multipath (ECMP); these paths are equal and the lowest cost from a routing perspective.

Step 6 Click **Submit**.

Creating a Load Balancer Policy Using the CLI

Creating a Dynamic Load Balancer Policy Using the CLI

There are two dynamic load balancer modes: **dynamic-aggressive** and **dynamic-conservative**. The **dynamic-aggressive** mode enables a shorter flowlet timeout interval, and the **dynamic-conservative** mode enables a longer flowlet timeout interval. For more information about these commands, see the *Cisco APIC NX-OS Style CLI Command Reference*.

This section demonstrates how to configure a dynamic load balancer policy using the CLI.

Procedure

Step 1 To enable aggressive mode dynamic load balancing:

```
apicl# conf t
apicl# (config)# system dynamic-load-balance mode dynamic-aggressive
```

Step 2 To enable conservative mode dynamic load balancing:

```
apicl# conf t
apicl# (config)# system dynamic-load-balance mode dynamic-conservative
```

Creating a Dynamic Packet Prioritization Policy Using the CLI

This section demonstrates how to enable dynamic packet prioritization using the CLI. For more information about this command, see the *Cisco APIC NX-OS Style CLI Command Reference*.

Procedure

Enable dynamic packet prioritization:

```
apicl# conf t
apicl# (config)# system dynamic-load-balance mode packet-prioritization
```

Creating a GTP Load Balancer Policy Using the CLI

This section demonstrates how to create a GTP load balancer policy using the CLI. For more information about this command, see the *Cisco APIC NX-OS Style CLI Command Reference*.

Procedure

Enable dynamic packet prioritization:

```
apicl# conf t
apicl# (config)# ip load-sharing address source_destination gtpu
```

Creating a Load Balancer Policy Using the REST API

This section demonstrates how to enable a DLB, DPP, and a GTP load balancer policy. For a list of all possible property values, see the *Cisco APIC Management Information Model Reference*.

Procedure

To enable a DLB, DPP, and GTP load balancer policy:

```
https://apic-ip-address/api/mo/uni.xml
<polUni>
<fabricInst>
  <lbpPol name="default" hashGtp="yes" pri="on" dlbMode="aggressive">
  </lbpPol>
</fabricInst>
</polUni>
```

Enable a Time Precision Policy

This topic describes how to enable Precision Time Protocol (PTP), a time synchronization protocol for nodes distributed across a network. Its hardware timestamp feature provides greater accuracy than other time synchronization protocols such as Network Time Protocol (NTP).

PTP is a distributed protocol that specifies how real-time PTP clocks in the system synchronize with each other. These clocks are organized into a master-member synchronization hierarchy with the grandmaster clock, the clock at the top of the hierarchy, determining the reference time for the entire system. Synchronization is achieved by exchanging PTP timing messages, with the members using the timing information to adjust their clocks to the time of their master in the hierarchy. PTP operates within a logical scope called a PTP domain.

Prior to release 4.2(5), the PTP function was supported only within the Cisco ACI fabric. Beginning in release 4.2(5), the PTP function is now supported outside of the Cisco ACI fabric, where a grandmaster and clients outside of the Cisco ACI fabric can be connected to front-panel ports on switches in the Cisco ACI fabric.



Note You can connect only one external grandmaster switch to a switch inside the Cisco ACI fabric that has been configured as a client to that grandmaster switch. Connecting multiple external grandmaster switches to a single client switch is not supported at this time.

Procedure

Step 1 On the menu bar, click **System > System Settings**.

Step 2 Click **Precision Time Protocol**.

Step 3 Choose **Enabled** or **Disabled**.

If you choose disable PTP, NTP time is used to sync the fabric. If you enable PTP, a spine is automatically chosen as a master to which the entire site gets synced.

Step 4 Click **Submit**.

Enable a Global System GIPo Policy

This topic describes how to use the infra tenant GIPo as the system GIPo.

An ACI multipod deployment requires the 239.255.255.240 system Global IP Outside (GIPo) to be configured on the inter-pod network (IPN) as a PIM BIDIR range. This 239.255.255.240 PIM BIDIR range configuration on the IPN devices can be avoided by using the infra GIPo as System GIPo.

Before you begin

Upgrade all of the switches in the ACI fabric, including the leaf switches and spine switches, to the latest APIC release.

Procedure

-
- Step 1** On the menu bar, click **System > System Settings**.
 - Step 2** Choose **Enabled** or **Disabled** (the default) on **Use Infra GIPo as System GIPo**
 - Step 3** Click **Submit**.
-

Configure a Fabric Port Tracking Policy

Uplink failure detection can be enabled in the fabric access fabric port tracking policy. The port tracking policy monitors the status of links between leaf switches and spine switches. When an enabled port tracking policy is triggered, the leaf switches take down all access interfaces on the switch that have EPGs deployed on them. For more information about fabric port tracking, see the *Cisco APIC Layer 2 Networking Configuration Guide*.

Procedure

-
- Step 1** On the menu bar, click **System > System Settings**.
 - Step 2** On the Navigation pane, choose **Port Tracking**.
 - Step 3** Enable port tracking by setting the **Port tracking state** to **on**.
 - Step 4** (Optional) Change the **Daily restore timer** value.
 - Step 5** Configure the **Number of active spine links that triggers port tracking** parameter.
 - Step 6** Click **Submit**.
-

Provisioning Global Fabric Access Policies

Create a Global Attachable Access Entity Profile

An Attachable Entity Profile (AEP) represents a group of external entities with similar infrastructure policy requirements. The infrastructure policies consist of physical interface policies that configure various protocol options, such as Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol (LLDP), or Link Aggregation Control Protocol (LACP).

An AEP is required to deploy VLAN pools on leaf switches. Encapsulation blocks (and associated VLANs) are reusable across leaf switches. An AEP implicitly provides the scope of the VLAN pool to the physical infrastructure.

The following AEP requirements and dependencies must be accounted for in various configuration scenarios, including network connectivity, VMM domains, and multipod configuration:

- The AEP defines the range of allowed VLANs but it does not provision them. No traffic flows unless an EPG is deployed on the port. Without defining a VLAN pool in an AEP, a VLAN is not enabled on the leaf port even if an EPG is provisioned.
- A particular VLAN is provisioned or enabled on the leaf port that is based on EPG events either statically binding on a leaf port or based on VM events from external controllers such as VMware vCenter or Microsoft Azure Service Center Virtual Machine Manager (SCVMM).
- Attached entity profiles can be associated directly with application EPGs, which deploy the associated application EPGs to all those ports associated with the attached entity profile. The AEP has a configurable generic function (infraGeneric), which contains a relation to an EPG (infraRsFuncToEpg) that is deployed on all interfaces that are part of the selectors that are associated with the attachable entity profile.

A virtual machine manager (VMM) domain automatically derives physical interface policies from the interface policy groups of an AEP.

Before you begin

Create the tenant, VRF instance, application profiles, and EPGs to associate to the attached entity profile.

Procedure

-
- | | |
|---------------|---|
| Step 1 | On the menu bar, click Fabric > Access Policies . |
| Step 2 | On the navigation bar, expand Policies and Global . |
| Step 3 | Right-click Attachable Access Entity Profile and choose Create Attachable Access Entity Profile . |
| Step 4 | Enter a name for the policy. |
| Step 5 | Click the + icon on Domains table. |
| Step 6 | Enter a physical domain, a previously created physical, Layer 2, Layer 3, or Fibre Channel domain, or create one. |
| Step 7 | Enter the encapsulation for the domain and click Update . |
| Step 8 | Click the + icon on the EPG DEPLOYMENT table. |

- Step 9** Enter the tenant, application profile, EPG, encapsulation (such as vlan-1), primary encapsulation (primary encapsulation number) and interface mode (trunk, Access (802.1P, or Access (Untagged).
 - Step 10** Click **Update**.
 - Step 11** Click **Next**.
 - Step 12** Choose the interfaces to associate to the attachable entity profile.
 - Step 13** Click **Finish**.
-

Configure the Global QoS Class Policy

The global QoS Class policy can be used to:

- Preserve the CoS priority level, to guarantee that the CoS value in 802.1P packets which enter and transit the ACI fabric is preserved. 802.1P CoS preservation is supported in single pod and multipod topologies. In multipod topologies, CoS Preservation can be used where you want to preserve the QoS priority settings of 802.1P traffic entering POD 1 and egressing out of POD 2, but you are not concerned with preserving the CoS/DSCP settings in interpod network (IPN) traffic between the pods. To preserve CoS/DSCP settings when multipod traffic is transiting an IPN, use a DSCP policy (configured at **Tenants > infra > > Policies > Protocol > DSCP class-cos translation policy for L3 traffic**)
- Reset the properties for the default QoS class levels, such as the **MTU**, **Queue Limit**, or **Scheduling Algorithm**.

Procedure

- Step 1** On the menu bar, click **Fabric > Access Policies**.
 - Step 2** On the navigation bar, expand **Policies** and **Global**.
 - Step 3** Click **QOS Class**.
 - Step 4** To enable 802.1P CoS preservation, click the **Preserve COS** check box.
 - Step 5** To change the default settings for a QoS class, double-click on it. Enter the new settings and click **Submit**.
-

Create a Global DHCP Relay Policy

The global DHCP relay policy identifies a DHCP relay server that can be referenced by bridge domains in different tenants across the fabric. With tenant DHCP relay you must configure the DHCP relay policy in each tenant even if the provider EPG is in the common tenant. The global DHCP relay policy differs from a tenant DHCP relay policy in that it can be referenced from any tenant and does not require the configuration of a DHCP relay policy in each tenant.

Procedure

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

- Step 2** On the navigation bar, expand **Policies** and **Global**.
- Step 3** Right-click **DHCP Relay** and choose **Create DHCP Relay Policy**.
- Step 4** In the **Create DHCP Relay Policy** dialog box, perform the following actions:
- In the **Name** field, enter the DHCP relay profile name (DhcpRelayP).
This name can be up to 64 alphanumeric characters.
 - (Optional) In the **Description** field, enter a description of the DHCP relay policy, if necessary.
The description can be up to 128 alphanumeric characters.
 - Expand **Providers**.
The **Create DHCP Provider** dialog box appears.
 - In the **Create DHCP Provider** dialog box, in the **EPG Type** field, click the appropriate radio button depending upon where the DHCP server is connected.
The options for the EPG type that you choose varies, depending on the EPG type.
 - If choose **Application EPG** as the EPG type, the following options appear in the **Application EPG** area:
 - In the **Tenant** field, from the drop-down list, choose the tenant. (infra)
 - In the **Application Profile** field, from the drop-down list, choose the application. (access)
 - In the **EPG** field, from the drop-down list, choose the EPG. (default)
 - If choose **L2 External Network** as the EPG type, the following options appear in the **L2 External Network** area:
 - In the **Tenant** field, from the drop-down list, choose the tenant.
 - In the **L2 Out** field, from the drop-down list, choose the L2 Out.
 - In the **External Network** field, from the drop-down list, choose the external network.
 - If choose **L3 External Network** as the EPG type, the following options appear in the **L3 External Network** area:
 - In the **Tenant** field, from the drop-down list, choose the tenant.
 - In the **L3 Out** field, from the drop-down list, choose the L3 Out.
 - In the **External Network** field, from the drop-down list, choose the external network.
 - If choose **DN** as the EPG type, enter the distinguished name of the target endpoint group.
 - In the **DHCP Server Address** field, enter the IP address for the DHCP server.
 - In the **DHCP Server Preference** field, choose **None** or **Use Server VRF**.
The **DHCP Server Preference** field is available beginning with release 5.2(4). Using the value in this field, the leaf switch determines whether to route the DHCP relay packets from the client VRF or the server VRF. For more information, see [About the DHCP Server Preference Field, on page 132](#).
 - None**: This is the default option, which reflects the behavior prior to release 5.2(4). By choosing the **None** option, the switch will always route the DHCP relay packet from the client VRF. If used for

inter-VRF DHCP relay, a shared services contract is required to leak the server VRF network into the client VRF.

- **Use Server VRF:** This option reflects new behavior introduced in release 5.2(4). By choosing the **Use Server VRF** option, the switch routes the DHCP relay packets from the server VRF, regardless of whether there is a contract or there is no contract between the EPG where the DHCP client is and the EPG where the DHCP server is (or the Layer 3 external of the L3Out through which the DHCP server is reachable).

For inter-VRF configurations, when you choose the **Use Server VRF** option in the **DHCP Server Preference** field, the server subnet route is programmed in the server VRF on the client leaf switch for route lookup. The DHCP process on the client leaf switch then sends the DHCP relay packets via the server VRF. Because of this, the server VRF must be deployed with at least one IP address on all leaf switches where the client bridge domains are deployed.

g) Click **OK**.

You are returned to the **Create DHCP Relay Policy** window.

h) Click **Submit**.

The DHCP relay policy is created.

Enable a Global MCP Instance Policy

Enable a global Mis-Cabling Protocol (MCP) instance policy. In the current implementation, only one instance of MCP runs in the system.

Procedure

- Step 1** On the menu bar, click **Fabric > Access Policies**.
- Step 2** On the navigation bar, expand **Policies** and **Global**.
- Step 3** Click **MCP Instance Policy default**.
- Step 4** Change the **Admin State** to **Enabled**.
- Step 5** Set other properties as needed for your fabric.
- Step 6** Click **Submit**.

What to do next

Create an Error Disabled Recovery Policy

The error disabled recovery policy specifies the policy for re-enabling a port that was disabled due to one or more pre-defined error conditions.

Procedure

-
- Step 1** On the menu bar, click **Fabric > Access Policies**.
 - Step 2** On the navigation bar, expand **Policies** and **Global**.
 - Step 3** Click **Error Disabled Recovery Policy**.
 - Step 4** Double-click on an event to enable it for the recovery policy.
 - Step 5** Click the check box and click **Update**.
 - Step 6** Optional. Repeat steps 4 and 5 for more events.
 - Step 7** Optional. Reset the **Error disable recovery interval (sec)**.
 - Step 8** Click **Submit**.
-

Per Port Policies

About Per Port Policies

A per port policy is an implicit policy that you use to configure the interfaces of a leaf switch using the Cisco Application Policy Infrastructure Controller (APIC) GUI. A per port policy is simplified compared to the standard policy-based model, which is useful when you are still learning how to use the Cisco APIC. Due to this simplification, you cannot add new ports to an existing policy. Instead, you can only create new chunks of a policy per interface.

The per port policy pane uses the NX-OS CLI in the background to create implicit and explicit objects. For example, creating a new port channel creates an explicit port channel policy group as well as an implicit override. Making any changes to an explicit policy group will not apply to port until the implicit policy group is removed. We recommend that you do not mix using the CLI and GUI. Use the per port policy wizard to learn about the Cisco Application Centric Infrastructure (ACI) policy model and unconfigure the port from the same wizard when you want to move to advanced use cases for reusable policy configurations.

You can create a per port policy only from the following GUI location:

Fabric > Inventory > Pod-# > leaf-switch-name > Interface tab



Note **Interface tab** refers to the **Interface** tab in the work pane. This is not the **Interfaces** folder in the navigation pane.

Configuring a Per Port Policy Using the GUI

This procedure configures a per port policy using the Cisco Application Policy Infrastructure Controller (APIC) GUI.

Procedure

-
- Step 1** On the menu bar, choose **Fabric > Inventory**.
- Step 2** In the Navigation pane, choose *pod-# > leaf-switch-name*.
- Step 3** In the Work pane, choose the **Interface** tab.
- Step 4** In the **Mode** drop-down list, choose **Configuration**.
- Step 5** Click on one or more of the interface numbers to select those interfaces.
- The buttons just under the tabs of the Work pane become active for any of the components that you can configure for the selected interfaces.
- Step 6** Click the button for one of the components that you want to configure.
- The Work pane displays the properties for that component.
- Step 7** Set the properties as desired for the component.
- Step 8** Click **Submit**.
- Step 9** Configure any additional components for the selected interfaces, or select different interfaces and configure the components.
-

Validating a Per Port Policy Using the GUI

This procedure instructs you on how to validate a per port policy using the Cisco Application Policy Infrastructure Controller (APIC) GUI.

Before you begin

You must configure the Cisco APIC to show hidden policies. By default, the per port policies are hidden in the Cisco APIC.

Procedure

-
- Step 1** On the menu bar, choose **Fabric > Inventory**.
- Step 2** In the Navigation pane, choose *pod-# > leaf-switch-name*.
- Step 3** In the Work pane, choose the **Interface** tab.
- Step 4** In the **Mode** drop-down list, choose **Configuration**.
- Step 5** Click on an interface numbers to select that interface.
- The buttons just under the tabs of the Work pane become active for any of the components that you can configure for the selected interface.
- Step 6** Click the button for one of the components for which you want to view the properties.
- The Work pane displays the properties for that component.

- Step 7** Verify that the properties are set correctly, and change any values that are not correct for your desired configuration.
- Step 8** If you made any changes, click **Submit**. Otherwise, click **Cancel**.
-

Showing the Hidden Policies Using the GUI

By default, some policies, such as per port policies, are hidden in the Cisco Application Policy Infrastructure Controller (APIC). If you want to view these policies, you must configure the Cisco APIC to show hidden policies.

Procedure

- Step 1** In the upper right corner of the GUI, choose **Manage My Profile > Settings**.
The **Application Settings** dialog opens.
- Step 2** Put a check in the **Show Hidden Policies** box.
- Step 3** Click **OK**.
-

Creating a Mis-cabling Protocol Interface Policy Using the GUI (Optional)

The mis-cabling protocol (MCP) was designed to handle misconfigurations that Link Layer Discovery Protocol (LLDP) and Spanning Tree Protocol (STP) are unable to detect. MCP has a Layer 2 packet that it uses, and MCP disables ports that form a loop within the Fabric. Cisco Application Centric Infrastructure (ACI) fabric leaf switches do not participate in spanning tree protocol (STP) and act as hub with respect to STP. The MCP packet is sent, and if the fabric sees that the packet comes back in, then the fabric knows that there is a loop and the fabric takes action based on that event. Faults and events are generated when this happens. MCP can be enabled globally and per-interface. By default, MCP is disabled globally and is enabled on each port. For MCP to work, it must be enabled globally, regardless of the per-interface configuration.

The following procedure creates an MCP interface policy using the GUI.

Procedure

- Step 1** On the menu bar, choose **Fabric > Access Policies**.
- Step 2** In the Navigation pane, choose **Policies > Interface > MCP Interface**.
- Step 3** In the Work pane, choose **Actions > Create Mis-cabling Protocol Interface Policy**.
- Step 4** In the **Create Mis-cabling Protocol Interface Policy** dialog box, perform the following actions:
- Enter a name for the policy.
 - (Optional) Enter a description for the policy.

- c) For **Admin State**, choose **Enable** to enable the policy, or **Disable** to disable the policy.
- d) Choose the **Strict** or **Non-strict** operational mode for MCP.

The following additional fields are displayed when you select **Strict**.

- **Initial Delay Time (sec)**: Time for STP convergence on the external Layer 2 network. Default value is 0 (when STP is disabled on the Layer 2 network). When STP is enabled, based on scale/ topology, the initial delay time range is 45 to 60 seconds for STP to converge.
- **Transmission Frequency (sec, msec)**: Transmission frequency timer for MCP packets until grace period on each Layer 2 interface. Default value is 500 milliseconds.
- **Grace Period (sec, msec)**: Grace period time during which early loop detection takes place. The port aggressively transmits MCP packets which are used for loop detection. Default grace period value is 3 seconds.

Step 5 Click **Submit**.



CHAPTER 7

Basic User Tenant Configuration


This chapter contains the following sections:

- [Tenants, on page 179](#)
- [Routing Within the Tenant, on page 180](#)
- [Creating Tenants, VRFs, and Bridge Domains, on page 191](#)
- [Deploying EPGs, on page 193](#)
- [Microsegmented EPGs, on page 198](#)
- [Deploying Application Profiles and Contracts, on page 202](#)
- [Optimize Contract Performance, on page 211](#)
- [Contract and Subject Exceptions, on page 214](#)
- [Intra-EPG Contracts, on page 216](#)
- [EPG Contract Inheritance, on page 224](#)
- [Contract Preferred Groups, on page 228](#)
- [Contracts with Permit and Deny Rules, on page 232](#)

Tenants

A tenant (`fvTenant`) is a logical container for application policies that enable an administrator to exercise domain-based access control. A tenant represents a unit of isolation from a policy perspective, but it does not represent a private network. Tenants can represent a customer in a service provider setting, an organization or domain in an enterprise setting, or just a convenient grouping of policies. The following figure provides an overview of the tenant portion of the management information tree (MIT).

Figure 1: Tenants



Tenants can be isolated from one another or can share resources. The primary elements that the tenant contains are filters, contracts, outside networks, bridge domains, Virtual Routing and Forwarding (VRF) instances, and application profiles that contain endpoint groups (EPGs). Entities in the tenant inherit its policies. VRFs are also known as contexts; each VRF can be associated with multiple bridge domains.



Note In the APIC GUI under the tenant navigation path, a VRF (context) is called a private network.

Tenants are logical containers for application policies. The fabric can contain multiple tenants. You must configure a tenant before you can deploy any Layer 4 to Layer 7 services. The ACI fabric supports IPv4, IPv6, and dual-stack configurations for tenant networking.

Routing Within the Tenant


The Application Centric Infrastructure (ACI) fabric provides tenant default gateway functionality and routes between the fabric virtual extensible local area (VXLAN) networks. For each tenant, the fabric provides a virtual default gateway or Switched Virtual Interface (SVI) whenever a subnet is created on the APIC. This spans any switch that has a connected endpoint for that tenant subnet. Each ingress interface supports the default gateway interface and all of the ingress interfaces across the fabric share the same router IP address and MAC address for a given tenant subnet.

Layer 3 VNIDs Facilitate Transporting Inter-subnet Tenant Traffic

The ACI fabric provides tenant default gateway functionality that routes between the ACI fabric VXLAN networks. For each tenant, the fabric provides a virtual default gateway that spans all of the leaf switches assigned to the tenant. It does this at the ingress interface of the first leaf switch connected to the endpoint. Each ingress interface supports the default gateway interface. All of the ingress interfaces across the fabric share the same router IP address and MAC address for a given tenant subnet.

The ACI fabric decouples the tenant endpoint address, its identifier, from the location of the endpoint that is defined by its locator or VXLAN tunnel endpoint (VTEP) address. Forwarding within the fabric is between VTEPs. The following figure shows decoupled identity and location in ACI.

Figure 2: ACI Decouples Identity and Location



VXLAN uses VTEP devices to map tenant end devices to VXLAN segments and to perform VXLAN encapsulation and de-encapsulation. Each VTEP function has two interfaces:

- A switch interface on the local LAN segment to support local endpoint communication through bridging
- An IP interface to the transport IP network


The IP interface has a unique IP address that identifies the VTEP device on the transport IP network known as the infrastructure VLAN. The VTEP device uses this IP address to encapsulate Ethernet frames and transmit the encapsulated packets to the transport network through the IP interface. A VTEP device also discovers the remote VTEPs for its VXLAN segments and learns remote MAC Address-to-VTEP mappings through its IP interface.

The VTEP in ACI maps the internal tenant MAC or IP address to a location using a distributed mapping database. After the VTEP completes a lookup, the VTEP sends the original data packet encapsulated in VXLAN with the destination address of the VTEP on the destination leaf switch. The destination leaf switch de-encapsulates the packet and sends it to the receiving host. With this model, ACI uses a full mesh, single hop, loop-free topology without the need to use the spanning-tree protocol to prevent loops.

The VXLAN segments are independent of the underlying network topology; conversely, the underlying IP network between VTEPs is independent of the VXLAN overlay. It routes the encapsulated packets based on the outer IP address header, which has the initiating VTEP as the source IP address and the terminating VTEP as the destination IP address.

The following figure shows how routing within the tenant is done.

Figure 3: Layer 3 VNIDs Transport ACI Inter-subnet Tenant Traffic



For each tenant VRF in the fabric, ACI assigns a single L3 VNID. ACI transports traffic across the fabric according to the L3 VNID. At the egress leaf switch, ACI routes the packet from the L3 VNID to the VNID of the egress subnet.

Traffic arriving at the fabric ingress that is sent to the ACI fabric default gateway is routed into the Layer 3 VNID. This provides very efficient forwarding in the fabric for traffic routed within the tenant. For example, with this model, traffic between 2 VMs belonging to the same tenant, on the same physical host, but on different subnets, only needs to travel to the ingress switch interface before being routed (using the minimal path cost) to the correct destination.

To distribute external routes within the fabric, ACI route reflectors use multiprotocol BGP (MP-BGP). The fabric administrator provides the autonomous system (AS) number and specifies the spine switches that become route reflectors.



Note Cisco ACI does not support IP fragmentation. Therefore, when you configure Layer 3 Outside (L3Out) connections to external routers, or Multi-Pod connections through an Inter-Pod Network (IPN), it is recommended that the interface MTU is set appropriately on both ends of a link.

IGP Protocol Packets (EIGRP, OSPFv3) are constructed by components based on the Interface MTU size. In Cisco ACI, if the CPU MTU size is less than the Interface MTU size and if the constructed packet size is greater than the CPU MTU, then the packet is dropped by the kernel, especially in IPv6. To avoid such control packet drops always configure the same MTU values on both the control plane and on the interface.

On some platforms, such as Cisco ACI, Cisco NX-OS, and Cisco IOS, the configurable MTU value does not take into account the Ethernet headers (matching IP MTU, and excluding the 14-18 Ethernet header size), while other platforms, such as IOS-XR, include the Ethernet header in the configured MTU value. A configured value of 9000 results in a max IP packet size of 9000 bytes in Cisco ACI, Cisco NX-OS, and Cisco IOS, but results in a max IP packet size of 8986 bytes for an IOS-XR untagged interface.


For the appropriate MTU values for each platform, see the relevant configuration guides.

We highly recommend that you test the MTU using CLI-based commands. For example, on the Cisco NX-OS CLI, use a command such as `ping 1.1.1.1 df-bit packet-size 9000 source-interface ethernet 1/1`.

Router Peering and Route Distribution

As shown in the figure below, when the routing peer model is used, the leaf switch interface is statically configured to peer with the external router's routing protocol.

Figure 4: Router Peering




The routes that are learned through peering are sent to the spine switches. The spine switches act as route reflectors and distribute the external routes to all of the leaf switches that have interfaces that belong to the same tenant. These routes are longest prefix match (LPM) summarized addresses and are placed in the leaf switch's forwarding table with the VTEP IP address of the remote leaf switch where the external router is connected. WAN routes have no forwarding proxy. If the WAN routes do not fit in the leaf switch's forwarding table, the traffic is dropped. Because the external router is not the default gateway, packets from the tenant endpoints (EPs) are sent to the default gateway in the ACI fabric.

Bridged Interface to an External Router

As shown in the figure below, when the leaf switch interface is configured as a bridged interface, the default gateway for the tenant VNID is the external router.

Figure 5: Bridged External Router



The ACI fabric is unaware of the presence of the external router and the APIC statically assigns the leaf switch interface to its EPG.

Configuring Route Reflectors

ACI fabric route reflectors use multiprotocol BGP (MP-BGP) to distribute external routes within the fabric. To enable route reflectors in the ACI fabric, the fabric administrator must select the spine switches that will be the route reflectors, and provide the autonomous system (AS) number. It is recommended to configure at least two spine nodes per pod as MP-BGP route reflectors for redundancy.

After route reflectors are enabled in the ACI fabric, administrators can configure connectivity to external networks through leaf nodes using a component called Layer 3 Out (L3Out). A leaf node configured with an L3Out is called a border leaf. The border leaf exchanges routes with a connected external device via a routing protocol specified in the L3Out. You can also configure static routes via L3Outs.

After both L3Outs and spine route reflectors are deployed, border leaf nodes learn external routes via L3Outs, and those external routes are distributed to all leaf nodes in the fabric via spine MP-BGP route reflectors.

Check the *Verified Scalability Guide for Cisco APIC* for your release to find the maximum number of routes supported by a leaf.

Configuring External Connectivity Using a Layer 3 Out

This section provides a step-by-step configuration required for the ACI fabric to connect to an external routed network through L3Outs and MP-BGP route reflectors.

This example uses Open Shortest Path First (OSPF) as the routing protocol in an L3Out under the 'mgmt' tenant.

Configuring an MP-BGP Route Reflector Using the GUI

Procedure

-
- Step 1** On the menu bar, choose **System > System Settings**.
- Step 2** In the **Navigation** pane, right-click **BGP Route Reflector**, and click **Create Route Reflector Node**.
- Step 3** In the **Create Route Reflector Node** dialog box, from the **Spine Node** drop-down list, choose the appropriate spine node. Click **Submit**.
- Note**
Repeat the above steps to add additional spine nodes as required.
- The spine switch is marked as the route reflector node.
- Step 4** In the **BGP Route Reflector** properties area, in the **Autonomous System Number** field, choose the appropriate number. Click **Submit**.
- Note**
The autonomous system number must match the leaf connected router configuration if Border Gateway Protocol (BGP) is configured on the router. If you are using routes learned using static or Open Shortest Path First (OSPF), the autonomous system number value can be any valid value.
- Step 5** On the menu bar, choose **Fabric > Fabric Policies > Pods > Policy Groups**.
- Step 6** In the **Navigation** pane, expand and right-click **Policy Groups**, and click **Create Pod Policy Group**.
- Step 7** In the **Create Pod Policy Group** dialog box, in the **Name** field, enter the name of a pod policy group.
- Step 8** In the **BGP Route Reflector Policy** drop-down list, choose the appropriate policy (default). Click **Submit**. The BGP route reflector policy is associated with the route reflector pod policy group, and the BGP process is enabled on the leaf switches.
- Step 9** On the menu bar, choose **Fabric > Fabric Policies > Profiles > Pod Profile default > default**.
- Step 10** In the **Work** pane, from the **Fabric Policy Group** drop-down list, choose the pod policy that was created earlier. Click **Submit**.
The pod policy group is now applied to the fabric policy group.
-

Configuring an MP-BGP Route Reflector for the ACI Fabric

To distribute routes within the ACI fabric, an MP-BGP process must first be operating, and the spine switches must be configured as BGP route reflectors.

The following is an example of an MP-BGP route reflector configuration:



Note In this example, the BGP fabric ASN is 100. Spine switches 104 and 105 are chosen as MP-BGP route-reflectors.

```
apic1(config)# bgp-fabric
apic1(config-bgp-fabric)# asn 100
apic1(config-bgp-fabric)# route-reflector spine 104,105
```

Configuring an MP-BGP Route Reflector Using the REST API

Procedure

Step 1 Mark the spine switches as route reflectors.

Example:

POST <https://apic-ip-address/api/policymgr/mo/uni/fabric.xml>

```
<bgpInstPol name="default">
  <bgpAsP asn="1" />
  <bgpRRP>
    <bgpRRNodePEp id="<spine_id1>" />
    <bgpRRNodePEp id="<spine_id2>" />
  </bgpRRP>
</bgpInstPol>
```

Step 2 Set up the pod selector using the following post.

Example:

For the FuncP setup—

POST <https://apic-ip-address/api/policymgr/mo/uni.xml>

```
<fabricFuncP>
  <fabricPodPGrp name="bgpRRPodGrp">
    <fabricRsPodPGrpBGPRRP tnBgpInstPolName="default" />
  </fabricPodPGrp>
</fabricFuncP>
```

Example:

For the PodP setup—

POST <https://apic-ip-address/api/policymgr/mo/uni.xml>

```
<fabricPodP name="default">
  <fabricPodS name="default" type="ALL">
    <fabricRsPodPGrp tDn="uni/fabric/funcprof/podpgrp-bgpRRPodGrp"/>
  </fabricPodS>
</fabricPodP>
```

Verifying the MP-BGP Route Reflector Configuration

Procedure

-
- Step 1** Verify the configuration by performing the following actions:
- Use secure shell (SSH) to log in as an administrator to each leaf switch as required.
 - Enter the **show processes | grep bgp** command to verify the state is S.
If the state is NR (not running), the configuration was not successful.
- Step 2** Verify that the autonomous system number is configured in the spine switches by performing the following actions:
- Use the SSH to log in as an administrator to each spine switch as required.
 - Execute the following commands from the shell window
- Example:**
- ```
cd /mit/sys/bgp/inst
```
- Example:**
- ```
grep asn summary
```
- The configured autonomous system number must be displayed. If the autonomous system number value displays as 0, the configuration was not successful.
-

Creating an OSPF L3Out for Management Tenant Using the GUI

- You must verify that the router ID and the logical interface profile IP address are different and do not overlap.
- The following steps are for creating an OSPF L3Out for a management tenant. To create an OSPF L3Out for a tenant, you must choose a tenant and create a VRF for the tenant.
- For more details, see *Cisco APIC and Transit Routing*.

Procedure

-
- Step 1** On the menu bar, choose **Tenants > mgmt**.
- Step 2** In the **Navigation** pane, expand **Networking > L3Outs**.
- Step 3** Right-click **L3Outs**, and click **Create L3Out**.
The **Create L3Out** wizard appears.
- Step 4** In the **Identity** window in the **Create L3Out** wizard, perform the following actions:
- In the **Name** field, enter a name (RtdOut).
 - In the **VRF** field, from the drop-down list, choose the VRF (inb).

Note

This step associates the routed outside with the in-band VRF.

- c) From the **L3 Domain** drop-down list, choose the appropriate domain.
- d) Check the **OSPF** check box.
- e) In the **OSPF Area ID** field, enter an area ID.
- f) In the **OSPF Area Control** field, check the appropriate check box.
- g) In the **OSPF Area Type** field, choose the appropriate area type.
- h) In the **OSPF Area Cost** field, choose the appropriate value.
- i) Click **Next**.

The **Nodes and Interfaces** window appears.

Step 5 In the **Nodes and Interfaces** window, perform the following actions:

- a) Uncheck the **Use Defaults** box.

This allows you to edit the **Node Profile Name** field.

- b) In the **Node Profile Name** field, enter a name for the node profile. (borderLeaf).
- c) In the **Node ID** field, from the drop-down list, choose the first node. (leaf1).
- d) In the **Router ID** field, enter a unique router ID.
- e) In the **Loopback Address** field, use a different IP address or leave this field empty if you do not want to use the router ID for the loopback address.

Note

The **Loopback Address** field is automatically populated with the same entry that you provide in the **Router ID** field. This is the equivalent of the **Use Router ID for Loopback Address** option in previous builds. Use a different IP address or leave this field empty if you do not want to use the router ID for the loopback address.

- f) Enter the appropriate information in the **Interface**, **IP Address**, **Interface Profile Name** and **MTU** fields for this node, if necessary.
- g) In the **Nodes** field, click + icon to add a second set of fields for another node.

Note

You are adding a second node ID.

- h) In the **Node ID** field, from the drop-down list, choose the first node. (leaf1).
- i) In the **Router ID** field, enter a unique router ID.
- j) In the **Loopback Address** field, use a different IP address or leave this field empty if you do not want to use the router ID for the loopback address.

Note

The **Loopback Address** field is automatically populated with the same entry that you provide in the **Router ID** field. This is the equivalent of the **Use Router ID for Loopback Address** option in previous builds. Use a different IP address or leave this field empty if you do not want to use the router ID for the loopback address.

- k) Enter the appropriate information in the **Interface**, **IP Address**, **Interface Profile Name** and **MTU** fields for this node, if necessary.
- l) Click **Next**.

The **Protocols** window appears.

Step 6 In the **Protocols** window, in the **Policy** area, click **default**, then click **Next**.
The **External EPG** window appears.

Step 7 In the **External EPG** window, perform the following actions:

- a) In the **Name** field, enter a name for the external network (extMgmt).
- b) Uncheck the **Default EPG for all external networks** field.
The **Subnets** area appears.
- c) Click + to access the **Create Subnet** dialog box.
- d) In the **Create Subnet** dialog box, in the **IP address** field, enter an IP address and mask for the subnet.
- e) In the **Scope** field, check the desired check boxes. Click **OK**.
- f) In the **External EPG** dialog box, click **Finish**.

Note

In the **Work** pane, in the **L3Outs** area, the L3Out icon (RtdOut) is now displayed.

Creating an OSPF External Routed Network for a Tenant Using the NX-OS CLI

Configuring external routed network connectivity involves the following steps:

1. Create a VRF under Tenant.
2. Configure L3 networking configuration for the VRF on the border leaf switches, which are connected to the external routed network. This configuration includes interfaces, routing protocols (BGP, OSPF, EIGRP), protocol parameters, route-maps.
3. Configure policies by creating external-L3 EPGs under tenant and deploy these EPGs on the border leaf switches. External routed subnets on a VRF which share the same policy within the ACI fabric form one "External L3 EPG" or one "prefix EPG".

Configuration is realized in two modes:

- Tenant mode: VRF creation and external-L3 EPG configuration
- Leaf mode: L3 networking configuration and external-L3 EPG deployment

The following steps are for creating an OSPF external routed network for a tenant. To create an OSPF external routed network for a tenant, you must choose a tenant and then create a VRF for the tenant.



Note The examples in this section show how to provide external routed connectivity to the "web" epg in the "OnlineStore" application for tenant "exampleCorp".

Procedure

Step 1 Configure the VLAN domain.

Example:

```
apicl(config)# vlan-domain dom_exampleCorp
apicl(config-vlan)# vlan 5-1000
apicl(config-vlan)# exit
```

Step 2 Configure the tenant VRF and enable policy enforcement on the VRF.

Example:

```
apicl(config)# tenant exampleCorp
apicl(config-tenant)# vrf context
    exampleCorp_v1
apicl(config-tenant-vrf)# contract enforce
apicl(config-tenant-vrf)# exit
```

Step 3 Configure the tenant BD and mark the gateway IP as “public”. The entry "scope public" makes this gateway address available for advertisement through the routing protocol for external-L3 network.

Example:

```
apicl(config-tenant)# bridge-domain exampleCorp_b1
apicl(config-tenant-bd)# vrf member exampleCorp_v1
apicl(config-tenant-bd)# exit
apicl(config-tenant)# interface bridge-domain exampleCorp_b1
apicl(config-tenant-interface)# ip address 172.1.1.1/24 scope public
apicl(config-tenant-interface)# exit
```

Step 4 Configure the VRF on a leaf.

Example:

```
apicl(config)# leaf 101
apicl(config-leaf)# vrf context tenant exampleCorp vrf exampleCorp_v1
```

Step 5 Configure the OSPF area and add the route map.

Example:

```
apicl(config-leaf)# router ospf default
apicl(config-leaf-ospf)# vrf member tenant exampleCorp vrf exampleCorp_v1
apicl(config-leaf-ospf-vrf)# area 0.0.0.1 route-map map100 out
apicl(config-leaf-ospf-vrf)# exit
apicl(config-leaf-ospf)# exit
```

Step 6 Assign the VRF to the interface (sub-interface in this example) and enable the OSPF area.

Example:

Note

For the sub-interface configuration, the main interface (ethernet 1/11 in this example) must be converted to an L3 port through “no switchport” and assigned a vlan-domain (dom_exampleCorp in this example) that contains the encapsulation VLAN used by the sub-interface. In the sub-interface ethernet1/11.500, 500 is the encapsulation VLAN.

```
apicl(config-leaf)# interface ethernet 1/11
apicl(config-leaf-if)# no switchport
apicl(config-leaf-if)# vlan-domain member dom_exampleCorp
apicl(config-leaf-if)# exit
apicl(config-leaf)# interface ethernet 1/11.500
apicl(config-leaf-if)# vrf member tenant exampleCorp vrf exampleCorp_v1
```

```
apic1(config-leaf-if)# ip address 157.10.1.1/24
apic1(config-leaf-if)# ip router ospf default area 0.0.0.1
```

- Step 7** Configure the external-L3 EPG policy. This includes the subnet to match for identifying the external subnet and consuming the contract to connect with the epg "web".

Example:

```
apic1(config)# tenant t100
apic1(config-tenant)# external-l3 epg l3epg100
apic1(config-tenant-l3ext-epg)# vrf member v100
apic1(config-tenant-l3ext-epg)# match ip 145.10.1.0/24
apic1(config-tenant-l3ext-epg)# contract consumer web
apic1(config-tenant-l3ext-epg)# exit
apic1(config-tenant)#exit
```

- Step 8** Deploy the external-L3 EPG on the leaf switch.

Example:

```
apic1(config)# leaf 101
apic1(config-leaf)# vrf context tenant t100 vrf v100
apic1(config-leaf-vrf)# external-l3 epg l3epg100
```

Creating Tenants, VRFs, and Bridge Domains

Tenants Overview

- A tenant contains policies that enable qualified users domain-based access control. Qualified users can access privileges such as tenant administration and networking administration.
- A user requires read/write privileges for accessing and configuring policies in a domain. A tenant user can have specific privileges into one or more domains.
- In a multitenancy environment, a tenant provides group user access privileges so that resources are isolated from one another (such as for endpoint groups and networking). These privileges also enable different users to manage different tenants.

Tenant Creation

A tenant contains primary elements such as filters, contracts, bridge domains, and application profiles that you can create after you first create a tenant.

VRF and Bridge Domains

You can create and specify a VRF and a bridge domain for the tenant. The defined bridge domain element subnets reference a corresponding Layer 3 context.

For details about enabling IPv6 Neighbor Discovery see *IPv6 and Neighbor Discovery* in *Cisco APIC Layer 3 Networking Guide*.

Creating a Tenant, VRF, and Bridge Domain Using the GUI

If you have a public subnet when you configure the routed outside, you must associate the bridge domain with the outside configuration.

Procedure

-
- Step 1** On the menu bar, choose **Tenants > Add Tenant**.
- Step 2** In the **Create Tenant** dialog box, perform the following tasks:
- In the **Name** field, enter a name.
 - In the **Security Domains** section, click the + to open the **Create Security Domain** dialog box.
 - In the **Name** field, enter a name for the security domain, then click **Submit**.
 - In the **Create Tenant** dialog box, click **Update** for the security domain that you created.
 - Fill in the other fields as necessary.
 - Click **Submit**.
- The *tenant_name* > **Networking** screen displays.
- Step 3** In the **Work** pane, drag the **VRF** icon to the canvas to open the **Create VRF** dialog box, and perform the following tasks:
- In the **Name** field, enter a name.
 - Fill in the other fields as necessary.
 - Click **Submit** to complete the VRF instance configuration.
- Step 4** In the **Work** pane, drag the **Bridge Domain** icon to the canvas within the circle around the VRF instance to connect the two. In the **Create Bridge Domain** dialog box that displays, perform the following tasks:
- In the **Name** field, enter a name.
 - Fill in the other fields as necessary.
 - Click **Next**.
 - In the **Subnets** section, click the + to open the **Create Subnet** dialog box.
 - In the **Gateway IP** field, enter the IP address and subnet mask.
 - Fill in the other fields as necessary.
 - Click **OK**.
 - Back in the **Create Bridge Domain** dialog box, fill in the other fields as necessary.
 - Click **Next**.
 - Fill in the fields as necessary.
 - Click **OK** to complete bridge domain configuration.
- Step 5** In the **Work** pane, drag the **L3** icon to the canvas within the circle around the VRF instance to connect the two. In the **Create Routed Outside** dialog box that displays, perform the following tasks:
- In the **Name** field, enter a name.
 - In the **Nodes And Interfaces Protocol Profiles** section, click the + to open the **Create Node Profile** dialog box.
 - In the **Name** field, enter a name.
 - In the **Nodes** section, click the + to open the **Select Node** dialog box.
 - In the **Node ID** drop-down list, choose a node.
 - In the **Router ID** field, enter the router ID.

- g) In the **Static Routes** section, click the + to open the **Create Static Route** dialog box.
- h) In the **Prefix** field, enter the IPv4 or IPv6 address.
- i) In the **Next Hop Addresses** section, click the + to open the **Create Next Hop** dialog box.
- j) In the **Next Hop Address** field, enter the IPv4 or IPv6 address.
- k) In the **Preference** field, enter a number.
- l) Fill in the other fields as necessary.
- m) Click **OK**.
- n) In the **Create Static Route** dialog box, fill in the other fields as necessary.
- o) Click **OK**.
- p) In the **Select Node** dialog box, fill in the other fields as necessary.
- q) Click **OK**.
- r) In the **Create Node Profile** dialog box, fill in the other fields as necessary.
- s) Click **OK**.
- t) Put a check in the **BGP**, **OSPF**, or **EIGRP** check boxes if desired.
- u) Fill in the other fields as necessary.
- v) Click **Next**.
- w) Fill in the fields as necessary.
- x) Click **OK** to complete the Layer 3 configuration.

To confirm the Layer 3 configuration, in the **Navigation** pane, expand **Networking** > **VRFs**.

Deploying EPGs

Statically Deploying an EPG on a Specific Port

This topic provides a typical example of how to statically deploy an EPG on a specific port when using Cisco APIC.

Deploying an EPG on a Specific Node or Port Using the GUI

Before you begin

The tenant where you deploy the EPG is already created.

You can create an EPG on a specific node or a specific port on a node.

Procedure

-
- Step 1** Log in to the Cisco APIC.
 - Step 2** Choose **Tenants** > *tenant*.
 - Step 3** In the left navigation pane, expand *tenant*, **Application Profiles**, and the *application profile*.
 - Step 4** Right-click **Application EPGs** and choose **Create Application EPG**.
 - Step 5** In the **Create Application EPG STEP 1 > Identity** dialog box, complete the following steps:
 - a) In the **Name** field, enter a name for the EPG.

- b) From the **Bridge Domain** drop-down list, choose a bridge domain.
- c) Check the **Statically Link with Leaves/Paths** check box.

This check box allows you to specify on which port you want to deploy the EPG.

- d) Click **Next**.
- e) From the **Path** drop-down list, choose the static path to the destination EPG.

Step 6 In the **Create Application EPG STEP 2 > Leaves/Paths** dialog box, from the **Physical Domain** drop-down list, choose a physical domain.

Step 7 Complete one of the following sets of steps:

Option	Description
If you want to deploy the EPG on...	Then
A node	<ul style="list-style-type: none"> a. Expand the Leaves area. b. From the Node drop-down list, choose a node. c. In the Encap field, enter the appropriate VLAN. d. (Optional) From the Deployment Immediacy drop-down list, accept the default On Demand or choose Immediate. e. (Optional) From the Mode drop-down list, accept the default Trunk or choose another mode.
A port on the node	<ul style="list-style-type: none"> a. Expand the Paths area. b. From the Path drop-down list, choose the appropriate node and port. c. (Optional) In the Deployment Immediacy field drop-down list, accept the default On Demand or choose Immediate. d. (Optional) From the Mode drop-down list, accept the default Trunk or choose another mode. e. In the Port Encap field, enter the secondary VLAN to be deployed. f. (Optional) In the Primary Encap field, enter the primary VLAN to be deployed.

Step 8 Click **Update** and click **Finish**.

Step 9 In the left navigation pane, expand the EPG that you created.

Step 10 Complete one of the following actions:

- If you created the EPG on a node, click **Static Leafs**, and in the work pane view details of the static binding paths.
- If you created the EPG on a port of the node, click **Static Ports**, and in the work pane view details of the static binding paths.

Creating Domains, Attach Entity Profiles, and VLANs to Deploy an EPG on a Specific Port

This topic provides a typical example of how to create physical domains, Attach Entity Profiles (AEP), and VLANs that are mandatory to deploy an EPG on a specific port.

All endpoint groups (EPGs) require a domain. Interface policy groups must also be associated with Attach Entity Profile (AEP), and the AEP must be associated with a domain, if the AEP and EPG have to be in same domain. Based on the association of EPGs to domains and of interface policy groups to domains, the ports and VLANs that the EPG uses are validated. The following domain types associate with EPGs:

- Application EPGs
- Layer 3 external outside network instance EPGs
- Layer 2 external outside network instance EPGs
- Management EPGs for out-of-band and in-band access

The APIC checks if an EPG is associated with one or more of these types of domains. If the EPG is not associated, the system accepts the configuration but raises a fault. The deployed configuration may not function properly if the domain association is not valid. For example, if the VLAN encapsulation is not valid for use with the EPG, the deployed configuration may not function properly.



Note EPG association with the AEP without static binding does not work in a scenario when you configure the EPG as **Trunk** under the AEP with one end point under the same EPG supporting Tagging and the other end point in the same EPG does not support VLAN tagging. While associating AEP under the EPG, you can configure it as Trunk, Access (Tagged) or Access (Untagged).

Creating Domains, and VLANs to Deploy an EPG on a Specific Port Using the GUI

Before you begin

- The tenant where you deploy the EPG is already created.
- An EPG is statically deployed on a specific port.

Procedure

-
- Step 1** On the menu bar, click **Fabric > Access Policies**.
- Step 2** In the **Navigation** pane, click **Quick Start**.
- Step 3** In the **Work** pane, click **Configure an Interface, PC, and vPC**.
- Step 4** In the **Configure an Interface, PC, and vPC** dialog box, click the + icon to select switches and perform the following actions:
- a) From the **Switches** drop-down list, check the check box for the desired switch.
 - b) In the **Switch Profile Name** field, a switch name is automatically populated.

Note

Optionally, you can enter a modified name.

- c) Click the + icon to configure the switch interfaces.
- d) In the **Interface Type** field, click the **Individual** radio button.
- e) In the **Interfaces** field, enter the range of desired interfaces.
- f) In the **Interface Selector Name** field, an interface name is automatically populated.

Note

Optionally, you can enter a modified name.

- g) In the **Interface Policy Group** field, choose the **Create One** radio button.
- h) From the **Link Level Policy** drop-down list, choose the appropriate link level policy.

Note

Create additional policies as desired, otherwise the default policy settings are available.

- i) From the **Attached Device Type** field, choose the appropriate device type.
- j) In the **Domain** field, click the **Create One** radio button.
- k) In the **Domain Name** field, enter a domain name.
- l) In the **VLAN** field, click the **Create One** radio button.
- m) In the **VLAN Range** field, enter the desired VLAN range. Click **Save**, and click **Save** again.
- n) Click **Submit**.

Step 5

On the menu bar, click **Tenants**. In the **Navigation** pane, expand the appropriate *Tenant_name* > **Application Profiles** > **Application EPGs** > *EPG_name* and perform the following actions:

- a) Right-click **Domains (VMs and Bare-Metals)**, and click **Add Physical Domain Association**.
- b) In the **Add Physical Domain Association** dialog box, from the **Physical Domain Profile** drop-down list, choose the appropriate domain.
- c) Click **Submit**.

The AEP is associated with a specific port on a node and with a domain. The physical domain is associated with the VLAN pool and the Tenant is associated with this physical domain.

The switch profile and the interface profile are created. The policy group is created in the port block under the interface profile. The AEP is automatically created, and it is associated with the port block and with the domain. The domain is associated with the VLAN pool and the Tenant is associated with the domain.

Deploying an Application EPG through an AEP or Interface Policy Group to Multiple Ports

Through the APIC Advanced GUI and REST API, you can associate attached entity profiles directly with application EPGs. By doing so, you deploy the associated application EPGs to all those ports associated with the attached entity profile in a single configuration.

Through the APIC REST API or the NX-OS style CLI, you can deploy an application EPG to multiple ports through an Interface Policy Group.

Deploying an EPG through an AEP to Multiple Interfaces Using the APIC GUI

You can quickly associate an application with an attached entity profile to quickly deploy that EPG over all the ports associated with that attached entity profile.

Before you begin

- The target application EPG is created.
- The VLAN pools has been created containing the range of VLANs you wish to use for EPG Deployment on the AEP.
- The physical domain has been created and linked to the VLAN Pool and AEP.
- The target attached entity profile is created and is associated with the ports on which you want to deploy the application EPG.

Procedure

- Step 1** Navigate to the target attached entity profile.
- Open the page for the attached entity profile to use. In the GUI, click **Fabric > Access Policies > Policies > Global > Attachable Access Entity Profiles**.
 - Click the target attached entity profile to open its Attachable Access Entity Profile window.

- Step 2** Click the **Show Usage** button to view the leaf switches and interfaces associated with this attached entity profile.
- the application EPGs associated with this attached entity profile are deployed to all the ports on all the switches associated with this attached entity profile.

- Step 3** Use the **Application EPGs** table to associate the target application EPG with this attached entity profile. Click + to add an application EPG entry. Each entry contains the following fields:

Field	Action
Application EPGs	Use the drop down to choose the associated Tenant, Application Profile, and target application EPG.
Encap	Enter the name of the VLAN over which the target application EPG will communicate.
Primary Encap	If the application EPG requires a primary VLAN, enter the name of the primary VLAN.
Mode	Use the drop down to specify the mode in which data is transmitted: <ul style="list-style-type: none"> • Trunk -- Choose if traffic from the host is tagged with a VLAN ID. • Access -- Choose if traffic from the host is tagged with an 802.1p tag. • Access Untagged -- Choose if the traffic from the host is untagged.

- Step 4** Click **Submit**.

the application EPGs associated with this attached entity profile are deployed to all the ports on all the switches associated with this attached entity profile.

Microsegmented EPGs

Using Microsegmentation with Network-based Attributes on Bare Metal

You can use Cisco APIC to configure Microsegmentation with Cisco ACI to create a new, attribute-based EPG using a network-based attribute, a MAC address or one or more IP addresses. You can configure Microsegmentation with Cisco ACI using network-based attributes to isolate VMs or physical endpoints within a single base EPG or VMs or physical endpoints in different EPGs.

Using an IP-based Attribute

You can use an IP-based filter to isolate a single IP address, a subnet, or multiple of noncontiguous IP addresses in a single microsegment. You might want to isolate physical endpoints based on IP addresses as a quick and simple way to create a security zone, similar to using a firewall.

Using a MAC-based Attribute

You can use a MAC-based filter to isolate a single MAC address or multiple MAC addresses. You might want to do this if you have a server sending bad traffic into the network. By creating a microsegment with a MAC-based filter, you can isolate the server.

Configuring Network-based Microsegmented EPGs in a Bare-Metal environment Using the GUI

You can use Cisco APIC to configure microsegmentation to put physical endpoint devices that belong to different base EPGs or the same EPG into a new attribute-based EPG.

Procedure

- Step 1** Log into the Cisco APIC.
- Step 2** Choose **Tenants** and then choose the tenant within which you want to create a microsegment.
- Step 3** In the tenant navigation pane, expand the tenant folder, the **Application Profiles** folder, the *profile* folder, and the **Application EPGs** folder.
- Step 4** Take one of the following actions:
- If you want to put physical endpoint devices from the same base EPG into a new, attribute-based EPG, click the base EPG containing the physical endpoint devices.
 - If you want to put physical endpoint devices from different base EPGs into a new, attribute-based EPG, click one of the base EPG containing the physical endpoint devices.
- The properties for the base EPG appear in the work pane.
- Step 5** In the work pane, click the **Operational** tab at the top right of the screen.
- Step 6** Below the **Operational** tab, ensure that the **Client End-Points** tab is active. The work pane displays all the physical endpoints that belong to the base EPG.

- Step 7** Note the IP address or MAC address for the endpoint device or endpoint devices that you want to put into a new microsegment.
- Step 8** If you want to put endpoint devices from different base EPGs into a new attribute-based EPG, repeat Step 4 through Step 7 for each of the base EPGs.
- Step 9** In the tenant navigation pane, right-click the **uSeg EPGs** folder, and then choose **Create uSeg EPG**.
- Step 10** Complete the following series of steps to begin creation of an attribute-based EPG for one of the groups of endpoint devices:
- In the **Create uSeg EPG** dialog box, in the **Name** field, enter a name.
We recommend that you choose a name that indicates that the new attribute-based EPG is a microsegment.
 - In the intra-EPG isolation field, select **enforced** or **unenforced**.
If you select **enforced**, ACI prevents all communication between the endpoint devices within this uSeg EPG.
 - In the **Bridge Domain** area, choose a bridge domain from the drop-down list.
 - In the **uSeg Attributes** area, choose **IP Address Filter** or **MAC Address Filter** from the + drop-down list on the right side of the dialog box.

Step 11 Complete one of the following series of steps to configure the filter.

If you want to use...	Then...
An IP-based attribute	<ol style="list-style-type: none"> In the Create IP Attribute dialog box, in the Name field, enter a name. We recommend that you choose a name that reflects the filter's function. In the IP Address field, enter an IP address or a subnet with the appropriate subnet mask. Click OK. (Optional) Create a second IP Address filter by repeating Step 10 c through Step 11 c. You might want to do this to include discontinuous IP addresses in the microsegment. In the Create uSeg EPG dialog box, click Submit.
A MAC-based attribute	<ol style="list-style-type: none"> In the Create MAC Attribute dialog box, in the Name field, enter a name. We recommend that you choose a name that reflects the filter's function. In the MAC Address field, enter a MAC address. Click OK. In the Create uSeg EPG dialog box, click Submit.

- Step 12** Complete the following steps to associate the uSeg EPG with a physical domain.
- In the navigation pane, ensure that the uSeg EPG folder is open and then open the container for the microsegment that you just created.
 - Click the folder **Domains (VMs and Bare-Metals)**.
 - On the right side of the work pane, click **Actions** and then choose **Add Physical Domain Association** from the drop-down list.

- d) In the **Add Physical Domain Association** dialog box, choose a profile from the **Physical Domain Profile** drop-down list.
- e) In the **Deploy Immediacy** area, accept the default **On Demand**.
- f) In the **Resolution Immediacy** area, accept the default **Immediate**.
- g) Click **Submit**.

Step 13 Associate the uSeg EPG with the appropriate leaf switch.

- a) In the navigation pane, ensure the uSeg EPG folder is open then click **Static Leafs**.
- b) In the Static Leafs window, click **Actions > Statically Link with Node**
- c) In the Statically Link With Node dialog, select the leaf node and mode.
- d) Click **Submit**.

Step 14 Repeat Step 9 through Step 13 for any other network attribute-based EPGs that you want to create.

What to do next

Verify that the attribute-based EPG was created correctly.

If you configured an IP-based or MAC-based attribute, make sure that traffic is running on the end point devices that you put into the new microsegments.

IP Address-Based Microsegmented EPG as a Shared Resource

You can configure an IP address-based microsegmented EPG as a resource that can be accessed from both within and without the VRF on which it is located. The method of doing so is to configure an existing IP address-based microsegmented EPG with a subnet (assigned a unicast IP address) and enable that subnet for being advertised and shared by devices located on VRFs other than the one on which this EPG is native. Then you define an IP attribute with an option enabled that associates the EPG with the IP address of the shared subnet.

Configuring an IP-based Microsegmented EPG as a Shared Resource Using the GUI

You can configure a microsegmented EPG with an IP-Address with 32 bit mask as a shared service, accessible by clients outside of the VRF and the current fabric.

Before you begin

The following GUI description of configuring assumes the preconfiguration of an IP address-based microsegmented EPG configured whose subnet mask is /32.



Note

- For directions on configuring an IP address based EPG in a physical environment, see [Using Microsegmentation with Network-based Attributes on Bare Metal, on page 198](#)
- For directions on configuring an IP address based EPG in a virtual environment, see *Configuring Microsegmentation with Cisco ACI* in the *Cisco ACI Virtualization Guide*.

Procedure

-
- Step 1** Navigate to the target IP-address-based EPG.
- In the APIC GUI, click **Tenant > tenant_name > uSeg EPGs > uSeg_epg_name** to display the EPG's **Properties** dialog.
- Step 2** For the target EPG, configure an IP attribute to match the EPG subnet address.
- In the **Properties** dialog, locate the **uSeg Attributes** table, and click +. When prompted, choose **IP Address Filter** to display the **Create IP Attribute** dialog.
 - Enter a name in the Name field
 - Check the box for **Use FV Subnet**.
Enabling this option, indicates that the IP attribute value matches the IP address of a shared subnet.
 - Click **Submit**.
- Step 3** Create a shared subnet for the target EPG.
- With the folder for the target IP address-based uSeg EPG still open in the APIC navigation pane, right-click the **Subnets** folder and select **Create EPG Subnets**.
 - In the **Default Gateway** field, enter the IP address/mask of the IP address-based microsegmented EPG.
- Note**
- In all cases the subnet mask must be /32.
 - In the context of an IP address-based EPG, you are not actually entering the default address for a gateway, rather you are entering the IP address for the shared EPG subnet.
- Select **Treat as a virtual IP address**.
 - Under Scope select **Advertised Externally** and **Shared between VRFs**.
 - Click **Submit**.
-

Unconfiguring an IP-based Microsegmented EPG as a Shared Resource Using the GUI

When you unconfigure an IP address-Based microsegmented EPG as a shared service, you must remove the shared subnet and also disable the option to use that subnet as a shared resource.

Before you begin

Before you unconfigure an IP address-based microsegmented EPG as a shared service, you should know the following:

- Know which subnet is configured as a shared service address for the IP address-based microsegmented EPG.
- Know which IP attribute is configured with the **Use FV Subnet** option enabled.

Procedure

-
- Step 1** Remove subnet from the IP address-based microsegmented EPG.
- In the APIC GUI, click **Tenant** > **tenant_name** > **Application Profiles** > **epg_name** > **uSeg EPGs** > **uSeg EPGs** > **uSeg_epg_name**.
 - With the folder for the target IP address-based uSeg EPG still open in the APIC navigation pane, click the **Subnets** folder.
 - In the **Subnets** window, select the subnet that is advertised and shared with other VRFs and click **Actions** > **Delete**. then
 - Click **Yes** to confirm the deletion.
- Step 2** Disable the **Use FV Subnet** option.
- With the folder for the target IP address-based uSeg EPG still open in the APIC navigation pane, click the name of the micro-segmented EPG to display the to display the EPG's **Properties** dialog.
 - In the **Properties** dialog, locate the **uSeg Attributes** table, and locate the IP attribute item with the **Use FV Subnet** option enabled.
 - Double-click that item to display the **Edit IP Attribute** dialog.
 - In the **Edit IP Attribute** dialog, deselect the **Use FV Subnet** option.
 - Assign another IP address attribute in the IP Address field.
- Note**
This address must be a unicast address with a 32 bit mask (for example: 124.124.124.123/32).
- Click **Submit**.
-

Deploying Application Profiles and Contracts

Security Policy Enforcement

As traffic enters the leaf switch from the front panel interfaces, the packets are marked with the EPG of the source EPG. The leaf switch then performs a forwarding lookup on the packet destination IP address within the tenant space. A hit can result in any of the following scenarios:

- A unicast (/32) hit provides the EPG of the destination endpoint and either the local interface or the remote leaf switch VTEP IP address where the destination endpoint is present.
- A unicast hit of a subnet prefix (not /32) provides the EPG of the destination subnet prefix and either the local interface or the remote leaf switch VTEP IP address where the destination subnet prefix is present.
- A multicast hit provides the local interfaces of local receivers and the outer destination IP address to use in the VXLAN encapsulation across the fabric and the EPG of the multicast group.



Note Multicast and external router subnets always result in a hit on the ingress leaf switch. Security policy enforcement occurs as soon as the destination EPG is known by the ingress leaf switch.

A miss result in the forwarding table causes the packet to be sent to the forwarding proxy in the spine switch. The forwarding proxy then performs a forwarding table lookup. If it is a miss, the packet is dropped. If it is a hit, the packet is sent to the egress leaf switch that contains the destination endpoint. Because the egress leaf switch knows the EPG of the destination, it performs the security policy enforcement. The egress leaf switch must also know the EPG of the packet source. The fabric header enables this process because it carries the EPG from the ingress leaf switch to the egress leaf switch. The spine switch preserves the original EPG in the packet when it performs the forwarding proxy function.

On the egress leaf switch, the source IP address, source VTEP, and source EPG information are stored in the local forwarding table through learning. Because most flows are bidirectional, a return packet populates the forwarding table on both sides of the flow, which enables the traffic to be ingress filtered in both directions.

Contracts Contain Security Policy Specifications

In the ACI security model, contracts contain the policies that govern the communication between EPGs. The contract specifies what can be communicated and the EPGs specify the source and destination of the communications. Contracts link EPGs, as shown below.

EPG 1 ----- CONTRACT ----- EPG 2

Endpoints in EPG 1 can communicate with endpoints in EPG 2 and vice versa if the contract allows it. This policy construct is very flexible. There can be many contracts between EPG 1 and EPG 2, there can be more than two EPGs that use a contract, and contracts can be reused across multiple sets of EPGs, and more.


There is also directionality in the relationship between EPGs and contracts. EPGs can either provide or consume a contract. An EPG that provides a contract is typically a set of endpoints that provide a service to a set of client devices. The protocols used by that service are defined in the contract. An EPG that consumes a contract is typically a set of endpoints that are clients of that service. When the client endpoint (consumer) tries to connect to a server endpoint (provider), the contract checks to see if that connection is allowed. Unless otherwise specified, that contract would not allow a server to initiate a connection to a client. However, another contract between the EPGs could easily allow a connection in that direction.

This providing/consuming relationship is typically shown graphically with arrows between the EPGs and the contract. Note the direction of the arrows shown below.

EPG 1 <-----consumes----- CONTRACT <-----provides----- EPG 2

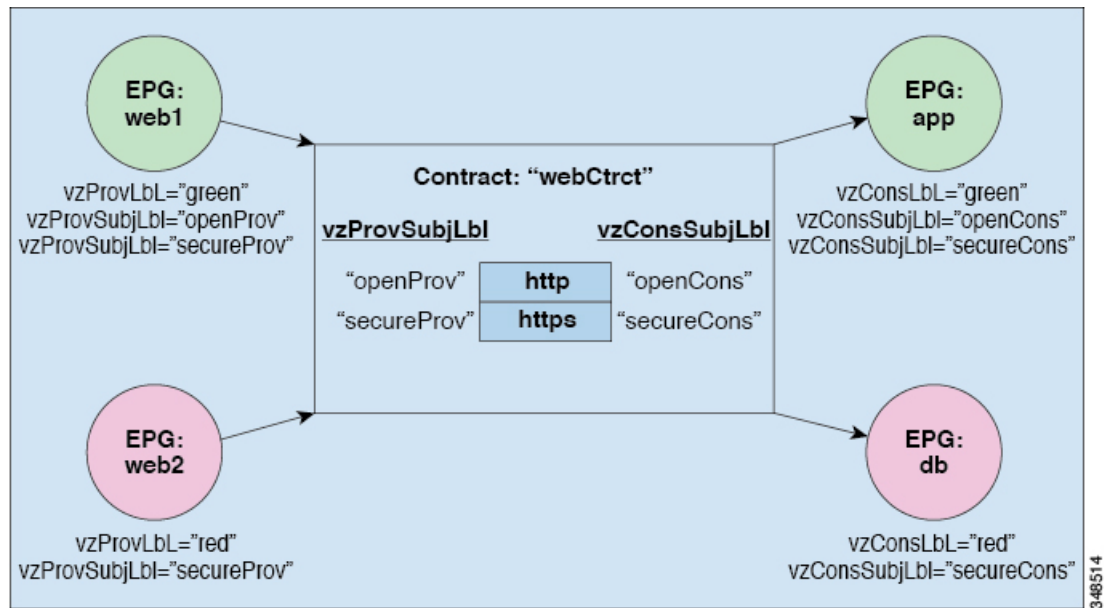
The contract is constructed in a hierarchical manner. It consists of one or more subjects, each subject contains one or more filters, and each filter can define one or more protocols.

Figure 6: Contract Filters



The following figure shows how contracts govern EPG communications.

Figure 7: Contracts Determine EPG to EPG Communications



For example, you may define a filter called HTTP that specifies TCP port 80 and port 8080 and another filter called HTTPS that specifies TCP port 443. You might then create a contract called webCtrct that has two sets of subjects. openProv and openCons are the subjects that contain the HTTP filter. secureProv and secureCons are the subjects that contain the HTTPS filter. This webCtrct contract can be used to allow both secure and non-secure web traffic between EPGs that provide the web service and EPGs that contain endpoints that want to consume that service.

These same constructs also apply for policies that govern virtual machine hypervisors. When an EPG is placed in a virtual machine manager (VMM) domain, the APIC downloads all of the policies that are associated with the EPG to the leaf switches with interfaces connecting to the VMM domain. For a full explanation of VMM domains, see the *Virtual Machine Manager Domains* chapter of *Application Centric Infrastructure Fundamentals*. When this policy is created, the APIC pushes it (pre-populates it) to a VMM domain that specifies which switches allow connectivity for the endpoints in the EPGs. The VMM domain defines the set

of switches and ports that allow endpoints in an EPG to connect to. When an endpoint comes on-line, it is associated with the appropriate EPGs. When it sends a packet, the source EPG and destination EPG are derived from the packet and the policy defined by the corresponding contract is checked to see if the packet is allowed. If yes, the packet is forwarded. If no, the packet is dropped.

Contracts consist of 1 or more subjects. Each subject contains 1 or more filters. Each filter contains 1 or more entries. Each entry is equivalent to a line in an Access Control List (ACL) that is applied on the Leaf switch to which the endpoint within the endpoint group is attached.

In detail, contracts are comprised of the following items:

- **Name**—All contracts that are consumed by a tenant must have different names (including contracts created under the common tenant or the tenant itself).
- **Subjects**—A group of filters for a specific application or service.
- **Filters**—Used to classify traffic based upon layer 2 to layer 4 attributes (such as Ethernet type, protocol type, TCP flags and ports).
- **Actions**—Action to be taken on the filtered traffic. The following actions are supported:
 - Permit the traffic (regular contracts, only)
 - Mark the traffic (DSCP/CoS) (regular contracts, only)
 - Redirect the traffic (regular contracts, only, through a service graph)
 - Copy the traffic (regular contracts, only, through a service graph or SPAN)
 - Block the traffic (taboo contracts)

With Cisco APIC Release 3.2(x) and switches with names that end in EX or FX, you can alternatively use a subject Deny action or Contract or Subject Exception in a standard contract to block traffic with specified patterns.
 - Log the traffic (taboo contracts and regular contracts)
- **Aliases**—(Optional) A changeable name for an object. Although the name of an object, once created, cannot be changed, the Alias is a property that can be changed.

Thus, the contract allows more complex actions than just allow or deny. The contract can specify that traffic that matches a given subject can be re-directed to a service, can be copied, or can have its QoS level modified. With pre-population of the access policy in the concrete model, endpoints can move, new ones can come on-line, and communication can occur even if the APIC is off-line or otherwise inaccessible. The APIC is removed from being a single point of failure for the network. Upon packet ingress to the ACI fabric, security policies are enforced by the concrete model running in the switch.

Three-Tier Application Deployment

A filter specifies the data protocols to be allowed or denied by a contract that contains the filter. A contract can contain multiple subjects. A subject can be used to realize uni- or bidirectional filters. A unidirectional filter is a filter that is used in one direction, either from consumer-to-provider (IN) or from provider-to-consumer (OUT) filter. A bidirectional filter is the same filter that is used in both directions. It is not reflexive.

Contracts are policies that enable inter-End Point Group (inter-EPG) communication. These policies are the rules that specify communication between application tiers. If no contract is attached to the EPG, inter-EPG


communication is disabled by default. No contract is required for intra-EPG communication because intra-EPG communication is always allowed.

Application profiles enable you to model application requirements that the APIC then automatically renders in the network and data center infrastructure. The application profiles enable administrators to approach the resource pool in terms of applications rather than infrastructure building blocks. The application profile is a container that holds EPGs that are logically related to one another. EPGs can communicate with other EPGs in the same application profile and with EPGs in other application profiles.

To deploy an application policy, you must create the required application profiles, filters, and contracts. Typically, the APIC fabric hosts a three-tier application within a tenant network. In this example, the application is implemented by using three servers (a web server, an application server, and a database server). See the following figure for an example of a three-tier application.

The web server has the HTTP filter, the application server has the Remote Method Invocation (RMI) filter, and the database server has the Structured Query Language (SQL) filter. The application server consumes the SQL contract to communicate with the database server. The web server consumes the RMI contract to communicate with the application server. The traffic enters from the web server and communicates with the application server. The application server then communicates with the database server, and the traffic can also communicate externally.

Figure 8: Three-Tier Application Diagram



Parameters to Create a Filter for http

The parameters to create a filter for http in this example is as follows:

Parameter Name	Filter for http
Name	http
Number of Entries	2
Entry Name	Dport-80 Dport-443
Ethertype	IP
Protocol	tcp tcp

Parameter Name	Filter for http
Destination Port	http https

Parameters to Create Filters for rmi and sql

The parameters to create filters for rmi and sql in this example are as follows:

Parameter Name	Filter for rmi	Filter for sql
Name	rmi	sql
Number of Entries	1	1
Entry Name	Dport-1099	Dport-1521
Ethertype	IP	IP
Protocol	tcp	tcp
Destination Port	1099	1521

Example Application Profile Database

The application profile database in this example is as follows:

EPG	Provided Contracts	Consumed Contracts
web	web	rmi
app	rmi	sql
db	sql	--

Creating an Application Profile Using the GUI

Procedure

-
- Step 1** On the menu bar, choose **TENANTS**. In the **Navigation** pane, expand the tenant, right-click **Application Profiles**, and click **Create Application Profile**.
- Step 2** In the **Create Application Profile** dialog box, in the **Name** field, add the application profile name (OnlineStore).
-

Creating EPGs Using the GUI

The port the EPG uses must belong to one of the VM Managers (VMM) or physical domains associated with the EPG.

Procedure

-
- Step 1** On the menu bar, choose **Tenants** and the tenant where you want to create an EPG.
- Step 2** In the navigation pane, expand the folder for the tenant, the **Application Profiles** folder, and the folder for the application profile.
- Step 3** Right-click the **Application EPG** folder, and in the **Create Application EPG** dialog box, perform the following actions:
- In the **Name** field, add the EPG name (db).
 - In the **Bridge Domain** field, choose the bridge domain from the drop-down list (bd1).
 - Check the **Associate to VM Domain Profiles** check box. Click **Next**.
 - In the **STEP 2 > Domains** area, expand **Associate VM Domain Profiles** and from the drop-down list, choose the desired VMM domain.
 - From the **Deployment Immediacy** drop-down list, accept the default or choose when policies are deployed from Cisco APIC to the physical leaf switch.
 - From the **Resolution Immediacy** drop-down list, choose when policies are deployed from the physical leaf switch to the virtual leaf.

If you have Cisco AVS, choose **Immediate** or **On Demand**; if you have Cisco ACI Virtual Edge or VMware VDS, choose **Immediate**, **On Demand**, or **Pre-provision**.
 - (Optional) In the **Delimiter** field, enter one of the following symbols: |, ~, !, @, ^, +, or =.

If you do not enter a symbol, the system uses the default | delimiter in the VMware portgroup name.
 - If you have Cisco ACI Virtual Edge or Cisco AVS, from the **Encap Mode** drop-down list, choose an encapsulation mode.

You can choose one of the following encapsulation modes:

 - VXLAN**: This overrides the domain's VLAN configuration, and the EPG uses VXLAN encapsulation. However, a fault is for the EPG if a multicast pool is not configured on the domain.
 - VLAN**: This overrides the domain's VXLAN configuration, and the EPG uses VLAN encapsulation. However, a fault is triggered for the EPG if a VLAN pool is not configured on the domain.
 - Auto**: This causes the EPG to use the same encapsulation mode as the VMM domain. This is the default configuration.
 - If you have Cisco ACI Virtual Edge, from the **Switching Mode** drop-down list, choose **native** or **AVE**.

If you choose **native**, the EPG is switched through the VMware VDS; if you choose **AVE**, the EPG is switched through the Cisco ACI Virtual Edge. The default is **native**.
 - Click **Update** and then click **Finish**.

- Step 4** In the **Create Application Profile** dialog box, create two more EPGs. Create the three EPGs—db, app, and web—in the same bridge domain and data center.
-

Configuring Contracts Using the APIC GUI

Guidelines and Limitations for Contracts and Filters

If your fabric consists of first-generation Cisco Nexus 9300 leaf switches, such as Cisco Nexus 93128TX, 93120TX, 9396TX, 9396PX, 9372PX, 9372PX-E, 9372TX and 9372TX-E, only **IP** as an **EtherType** match is supported with contract filters. The capability to match more granular options, such as **IPv4** or **IPv6**, in the **EtherType** field for contract filters is supported only on leaf switch models with -EX, -FX, or -FX2 at the end of the switch name.

Creating a Filter Using the GUI

Create three separate filters. In this example they are HTTP, RMI, SQL. This task shows how to create the HTTP filter. The task is identical for creating the other filters.

Before you begin

Verify that the tenant, network, and bridge domain have been created.

Procedure

- Step 1** On the menu bar, choose **Tenants**.
- Step 2** In the **Navigation** pane, expand the *tenant-name* > **Contracts**, right-click **Filters**, and choose **Create Filter**.
- Step 3** In the **Create Filter** dialog box, perform the following actions:
- In the **Name** field, enter the filter name (http).
 - In the **Entries** table, click +.and in the **Name** field, enter the name (Dport-80).
 - From the **EtherType** drop-down list, choose the EtherType (IP).
 - From the **IP Protocol** drop-down list, choose the protocol (tcp).
 - From the **Destination Port/Range** drop-down lists, choose **http** in the **From** and **To** fields. (http)
 - Click **Update**, and click **Submit**.

The newly added filter appears in the **Navigation** pane and in the **Work** pane.

Note

In the **Entries** table, **ARP Flag** has no functionality and you cannot configure it. Ignore this field.

- Step 4** Expand **Entries** in the **Name** field. Follow the same process to add another entry with HTTPS as the **Destination** port, and click **Update**.

This new filter rule is added.

- Step 5** Follow the same process in the earlier steps to create two more filters (rmi and sql) and use the parameters provided in [Parameters to Create Filters for rmi and sql](#), on page 207.
-

Creating a Contract Using the GUI

Procedure

-
- Step 1** On the menu bar, choose **Tenants** and the tenant name on which you want to operate. In the **Navigation** pane, expand the *tenant-name* > **Contracts**.
- Step 2** Right-click **Standard** > **Create Contract**.
- Step 3** In the **Create Contract** dialog box, perform the following tasks:
- In the **Name** field, enter the contract name (web).
 - Click the + sign next to **Subjects** to add a new subject.
 - In the **Create Contract Subject** dialog box, enter a subject name in the **Name** field. (web)
 - Note**
This step associates the filters created that were earlier with the contract subject.
- In the **Filter Chain** area, click the + sign next to **Filters**.
- In the dialog box, from the drop-down menu, choose the filter name (http), and click **Update**.
- Step 4** In the **Create Contract Subject** dialog box, click **OK**.
- Step 5** Create two more contracts for rmi and for sql following the same steps in this procedure. For the rmi contract, choose the rmi subject and for sql, choose the sql subject.
-

Consuming and Providing Contracts Using the GUI

You can associate contracts that were created earlier to create policy relationships between the EPGs.

When you name the provided and consumed contracts, verify that you give the same name for both provided and consumed contracts.

Procedure

-
- Step 1** **Note**
The db, app, and web EPGs are displayed as icons.
- Click and drag across the APIC GUI window from the db EPG to the app EPG.
The **Add Consumed Contract** dialog box is displayed.
- Step 2** In the **Name** field, from the drop-down list, choose **sql** contract. Click **OK**.
This step enables the db EPG to provide the sql contract and the app EPG to consume the sql contract.
- Step 3** Click and drag across the APIC GUI screen from the app ePG to the web EPG.
The **Add Consumed Contract** dialog box is displayed.
- Step 4** In the **Name** field, from the drop-down list, choose **rmi** contract. Click **OK**.
This step enables the app EPG to provide the rmi contract and the web EPG to consume the rmi contract.
- Step 5** Click the web EPG icon, and click the + sign in the **Provided Contracts** area.
The **Add Provided Contract** dialog box is displayed.
- Step 6** In the **Name** field, from the drop-down list, choose **web** contract. Click **OK**. Click **Submit**.
You have created a three-tier application profile called OnlineStore.

- Step 7** To verify, in the **Navigation** pane, navigate to and click **OnlineStore** under **Application Profiles**. In the **Work** pane, you can see the three EPGs app, db, and web are displayed.
- Step 8** In the **Work** pane, choose **Operational > Contracts**. You can see the EPGs and contracts displayed in the order that they are consumed and provided.
-

Optimize Contract Performance

Optimize Contract Performance

Starting with Cisco APIC, Release 3.2, you can configure bidirectional contracts that support more efficient hardware TCAM storage of contract data. With optimization enabled, contract statistics for both directions are aggregated.

TCAM Optimization is supported on the second generation Cisco Nexus 9000 Series top of rack (TOR) switches, which are those with suffixes of EX, FX, and FX2, and later (for example, N9K-C93180LC-EX or N9K-C93180YC-FX).

To configure efficient TCAM contract data storage, you enable the following options:

- Mark the contracts to be applied in both directions between the provider and consumer.
- For filters with IP TCP or UDP protocols, enable the reverse port option.
- When configuring the contract subjects, select the **Enable Policy Compression** directive, which adds the `no_stats` option to the `action` attribute of the `actrl:Rule` managed object.

Limitations

With the **Enable Policy Compression** (`no_stats`) option selected, per-rule statistics are lost. However, combined rule statistics for both directions are present in the hardware statistics.

After upgrading to Cisco APIC 3.2(1), to add the `no_stats` option to a pre-upgrade contract subject (with filters or filter entries), you must delete the contract subject and reconfigure it with the **Enable Policy Compression** directive. Otherwise, compression does not occur.

For each contract with a bi-directional subject filter, Cisco NX-OS creates 2 rules:

- A rule with an `sPcTag` and `dPcTag` that is marked `direction=bi-dir`, which is programmed in hardware
- A rule marked with `direction=uni-dir-ignore` which is not programmed

Rules with the following settings are not compressed:

- Rules with priority other than `fully_qual`
- Opposite rules (`bi-dir` and `uni-dir-ignore` marked) with non-identical properties, such as **action** including **directives**, **prio**, **qos** or **markDscp**
- Rule with `Implicit` or `implarp` filters
- Rules with the actions `Deny`, `Redir`, `Copy`, or `Deny-log`

The following MO query output shows the two rules for a contract, that is considered for compression:

```

apic1# moquery -c actrlRule
Total Objects shown: 2

# actrl.Rule
scopeId          : 2588677
sPcTag           : 16388
dPcTag           : 49156
fltId            : 67
action           : no_stats,permit
actrlCfgFailedBmp :
actrlCfgFailedTs : 00:00:00:00.000
actrlCfgState    : 0
childAction      :
ctrctName        :
descr            :
direction        : bi-dir
dn               : sys/actrl/scope-2588677/rule-2588677-s-16388-d-49156-f-67
id               : 4112
lcOwn            : implicit
markDscp        : unspecified
modTs           : 2019-04-27T09:01:33.152-07:00
monPolDn        : uni/tn-common/monepg-default
name             :
nameAlias        :
operSt           : enabled
operStQual       :
prio             : fully_qual
qosGrp          : unspecified
rn               : rule-2588677-s-16388-d-49156-f-67
status           :
type             : tenant

# actrl.Rule
scopeId          : 2588677
sPcTag           : 49156
dPcTag           : 16388
fltId            : 64
action           : no_stats,permit
actrlCfgFailedBmp :
actrlCfgFailedTs : 00:00:00:00.000
actrlCfgState    : 0
childAction      :
ctrctName        :
descr            :
direction        : uni-dir-ignore
dn               : sys/actrl/scope-2588677/rule-2588677-s-49156-d-16388-f-64
id               : 4126
lcOwn            : implicit
markDscp        : unspecified
modTs           : 2019-04-27T09:01:33.152-07:00
monPolDn        : uni/tn-common/monepg-default
name             :
nameAlias        :
operSt           : enabled
operStQual       :
prio             : fully_qual
qosGrp          : unspecified
rn               : rule-2588677-s-49156-d-16388-f-64
status           :
type             : tenant

```

Table 7: Compression Matrix

Reverse Filter Port Enabled	TCP or UDP Source Port	TCP or UCP Destination Port	Compressed
Yes	Port A	Port B	Yes
Yes	Unspecified	Port B	Yes
Yes	Port A	Unspecified	Yes
Yes	Unspecified	Unspecified	Yes
No	Port A	Port B	No
No	Unspecified	Port B	No
No	Port A	Unspecified	No
No	Unspecified	Unspecified	Yes

Configure a Contract with Optimized TCAM Usage Using the GUI

This procedure describes how to configure a contract that optimizes TCAM storage of contract data on hardware.

Before you begin

- Create the tenant, VRF, and EPGs that will provide and consume the contract.
- Create one or more filters that define the traffic to be permitted or denied by this contract.

Procedure

-
- Step 1** On the menu bar, choose **Tenants** and the tenant name on which you want to operate. In the **Navigation** pane, expand the *tenant-name* and **Contracts**.
- Step 2** Right-click **Standard** > **Create Contract**.
- Step 3** In the **Create Contract** dialog box, perform the following tasks:
- In the **Name** field, enter the contract name.
 - Click the + icon next to **Subjects** to add a new subject.
 - In the **Create Contract Subject** dialog box, enter a subject name in the **Name** field.

Note

This step associates filters with the contract subject.

- To enable the TCAM-contract usage optimization feature, ensure that **Apply Both Directions** and **Reverse Filter Ports** are enabled.
- Click the + icon to expand **Filters**.

- f) In the dialog box, from the drop-down menu, choose a default filter, a previously configured filter, or **Create Filter**.
- g) In the **Directives** field, choose **Enable Policy Compression**
- h) In the **Action** field, choose **Permit** or **Deny**.

Note

Currently, the **Deny** action is not supported. Optimization only occurs for the **Permit** action.

- i) (Optional) In the **Priority** field, choose the priority level.
- j) Click **Update**.

Step 4 In the **Create Contract Subject** dialog box, click **OK**.

Step 5 In the **Create Contract** dialog box, click **Submit**.

Contract and Subject Exceptions

Configuring Contract or Subject Exceptions for Contracts

In Cisco APIC Release 3.2(1), contracts between EPGs are enhanced to enable denying a subset of contract providers or consumers from participating in the contract. Inter-EPG contracts and Intra-EPG contracts are supported with this feature.

You can enable a provider EPG to communicate with all consumer EPGs except those that match criteria configured in a subject or contract exception. For example, if you want to enable an EPG to provide services to all EPGs for a tenant, except a subset, you can enable those EPGs to be excluded. To configure this, you create an exception in the contract or one of the subjects in the contract. The subset is then denied access to providing or consuming the contract.

Labels, counters, and permit and deny logs are supported with contracts and subject exceptions.

To apply an exception to all subjects in a contract, add the exception to the contract. To apply an exception only to a single subject in the contract, add the exception to the subject.

When adding filters to subjects, you can set the action of the filter (to permit or deny objects that match the filter criteria). Also for **Deny** filters, you can set the priority of the filter. **Permit** filters always have the default priority. Marking the subject-to-filter relation to deny automatically applies to each pair of EPGs where there is a match for the subject. Contracts and subjects can include multiple subject-to-filter relationships that can be independently set to permit or deny the objects that match the filters.

Exception Types

Contract and subject exceptions can be based on the following types and include regular expressions, such as the * wildcard:

Exception criteria exclude these objects as defined in the Consumer Regex and Provider Regex fields	Example	Description
Tenant	<pre><vzException consRegex= "common" field= "Tenant" name= "excep03" provRegex= "t1" /></pre>	This example, excludes EPGs using the <code>common</code> tenant from consuming contracts provided by the <code>t1</code> tenant.
VRF	<pre><vzException consRegex= "ctx1" field= "Ctx" name= "excep05" provRegex= "ctx1" /></pre>	This example excludes members of <code>ctx1</code> from consuming the services provided by the same VRF.
EPG	<pre><vzException consRegex= "EPgPa.*" field= "EPg" name= "excep03" provRegex= "EPg03" /></pre>	The example assumes that multiple EPGs exist, with names starting with <code>EPgPa</code> , and they should all be denied as consumers for the contract provided by <code>EPg03</code> .
Dn	<pre><vzException consRegex= "uni/tn-t36/ap-customer/epg-epg193" field= "Dn" name="excep04" provRegex= "uni/tn-t36/ap-customer/epg-epg200" /></pre>	This example excludes <code>epg193</code> from consuming the contract provided by <code>epg200</code> .
Tag	<pre><vzException consRegex= "red" field= "Tag" name= "excep01" provRegex= "green" /></pre>	The example excludes objects marked with the <code>red</code> tag from consuming and those marked with the <code>green</code> tag from participating in the contract.

Configure a Contract or Subject Exception Using the GUI

In this task, you configure a contract that will allow most of the EPGs to communicate, but deny access to a subset of them.

Before you begin

Configure the tenant, VRF, application profile, and EPGs that provide and consume the contract.

Procedure

-
- Step 1** Click **Tenants** > **All Tenants** on the menu bar.
- Step 2** Double-click the tenant in which you are creating the contract.
- Step 3** On the navigation bar, expand **Contracts**, right-click **Filter**, and choose **Create Filter**.
- A filter is essentially an Access Control List (ACL) that defines the traffic that is permitted or denied access through the contract. You can create multiple filters that define objects that can be permitted or denied.
- Step 4** Enter the filter name and add the criteria that define the traffic to permit or deny, then click **Submit**.

- Step 5** Right-click **Standard**, and choose **Create Contract**.
- Step 6** Enter the contract name, set the scope, and click the + icon to add a subject.
- Step 7** Repeat to add another subject.
- Step 8** Click **Submit**
- Step 9** To add an exception to all subjects in the contract, perform the following steps:
- Click the contract, then click **Contract Exception**.
 - Add subjects and set them to be permitted or denied.
 - Click the + icon to add a contract exception.
 - Enter the exception name and type.
 - Add regular expressions in the **Consumer Regex** and **Provider Regex** fields to define the EPGs to be excluded from all subjects in the contract.
- Step 10** To add an exception to one subject in the contract, perform the following steps:
- Click the subject, then click **Subject Exception**.
 - Click the + icon to add a contract exception.
 - Enter the exception name and type.
 - Add regular expressions in the **Consumer Regex** and **Provider Regex** to define the EPGs to be excluded from all subjects in the contract.
-

Intra-EPG Contracts

Intra-EPG Contracts

You can configure contracts to control communication between EPGs. Beginning in Cisco APIC Release 3.0(1), you can also configure contracts within an EPG.

Without intra-EPG contracts, communication between endpoints in an EPG is all-or-nothing. Communication is unrestricted by default, or you can configure intra-EPG isolation to bar any communication between endpoints.

However, with intra-EPG contracts, you can control communication between endpoints in the same EPG, allowing some traffic and barring the rest. For example, you may want to allow web traffic but block the rest. Or you can allow all ICMP traffic and TCP port 22 traffic while blocking all other traffic.

Guidelines and Limitations for Intra-EPG Contracts

Observe the following guidelines and limitations when planning intra-EPG contracts:

- Intra-EPG contracts can be configured for application EPGs and microsegment EPGs (uSegs) on VMware VDS, Open vSwitch (OVS), and baremetal servers.



Note OVS is available in the Kubernetes integration with Cisco Application Centric Infrastructure (ACI) feature. In Kubernetes, you can create EPGs and assign namespaces to them. You can then apply intra-EPG policies to the EPGs in Cisco Application Policy Infrastructure Controller (APIC) as you would for VMware VDS or baremetal servers.

- Intra-EPG contracts require that the leaf switch support proxy Address Resolution Protocol (ARP). Intra-EPG contracts are supported on Cisco Nexus 9000 Series switches with EX or FX at the end of their model name or later models.
- Intra-EPG Contracts are not supported in Cisco Application Virtual Switch, Cisco ACI Virtual Edge, and Microsoft domains. Attempting to set intra-EPG contracts to be enforced in these domains may cause ports to go into a blocked state.
- Intra-EPG contracts in service graphs:
 - A service graph cannot be associated with a subject of an intra-EPG contract that has an action of deny.
 - Support for intra-EPG contracts in service graphs is limited to single node one-arm mode policy-based redirect and copy service.
- Beginning in Cisco APIC release 5.2(1), intra-EPG contracts are supported on L3Out EPGs.
 - The action can be **permit**, **deny**, or **redirect**. The **redirect** action requires a service graph with policy-based redirect (PBR).
 - An L3Out EPG with an IP address and subnet of 0.0.0.0/0 or 0:::0 cannot use an intra-EPG contract nor intra-EPG isolation. The Cisco APIC raises a fault in these cases. However, you can instead use an IP address and subnet of 0.0.0.0/1 and 128.0.0.0/1 for the L3Out EPG to catch all traffic.
 - Unlike an intra-EPG contract on an EPG, the implicit deny rule is not automatically added for an intra-EPG contract on an L3Out EPG. You must enable intra-EPG isolation to deny other traffic. Intra-EPG isolation on an L3Out EPG works only when the VRF instance is in the enforced mode.
 - Cisco ACI cannot control how traffic reaches the Cisco ACI border leaf switch for intra L3Out enforcement.

Adding an Intra-EPG Contract to an Application EPG Using the GUI

After you configure a contract, you can add the contract to an EPG as an intra-EPG contract. The procedure is the same for VMware VDS, OVS, and baremetal servers.

Before you begin

- You must have an application EPG configured.
- You must have a contract with filters configured for this application. See [Creating a Contract Using the GUI, on page 210](#).

Procedure

Step 1 On the menu bar, choose **Tenants > All Tenants**.

Step 2 In the Work pane, double click the tenant's name.

Step 3 Complete one of the following sets of steps, depending on the type of EPG:

If you want to apply an intra-EPG contract to...	Then...
An application EPG	<ol style="list-style-type: none"> a. In the Navigation pane, expand <i>tenant_name</i> > Application Profiles > <i>application profile</i> > Application EPGs > <i>epg</i>. b. Right-click the Contracts folder and then choose Add Intra-EPG Contract. c. In the Add Intra Ext-EPG Contract dialog box, from the Contract drop-down list, choose an existing contract or create a new contract. d. Click Submit.
A uSeg EPG	<ol style="list-style-type: none"> a. In the Navigation pane, expand <i>tenant_name</i> > Application Profiles > <i>application profile</i> > uSeg EPGs > <i>epg</i>. b. Right-click the Contracts folder and then choose Add Intra-EPG Contract. c. In the Add Intra Ext-EPG Contract dialog box, from the Contract drop-down list, choose an existing contract or create a new contract. d. Click Submit.
An L3Out EPG	<ol style="list-style-type: none"> a. In the Navigation pane, choose <i>tenant_name</i> > Networking > L3Outs > <i>L3Out_name</i> > External EPGs > <i>ext_epg_name</i>. b. In the Work pane, for Intra Ext-EPG Isolation, choose Enforced. c. Click Submit. d. In the Work pane, choose the Policy > Contracts tab. e. In the action menu, choose Add Intra Ext-EPG Contract. f. In the Add Intra Ext-EPG Contract dialog box, from the Contract drop-down list, choose an existing contract or create a new contract. g. Click Submit. <p>The chosen contract appears in the Contract Type: Intra EPG Contract section of the Work pane.</p>

Adding an Intra-EPG Contract to an Application EPG Using the NX-OS-Style CLI

After you configure a contract, you can configure the contract as an intra-EPG contract. The procedure is the same for VMware VDS, OVS, and baremetal servers.

Before you begin

- You must have an EPG configured.
- You must have a contract with filters configured.

Procedure

Step 1 Enter the configuration mode.

Example:

```
apic1# configure
```

Step 2 Create or choose a tenant.

Example:

```
apic1(config)# tenant Tenant-13out
```

Step 3 Create or choose an external Layer 3 EPG.

Example:

```
apic1(config-tenant)# external-13 epg ext-epg
```

Step 4 Bind the external EPG to a VRF instance.

Example:

```
apic1(config-tenant-13ext-epg)# vrf member vrf1
```

Step 5 Enable intra-EPG isolation.

Example:

```
(config-tenant-13ext-epg)# isolation enforce
```

Afterward, if necessary, you can disable intra-EPG isolation by preceding the command with `no`.

Example:

```
(config-tenant-13ext-epg)# no isolation enforce
```

Step 6 Assign a contract to the intra-external EPG to allow the desired traffic between the endpoints.

Example:

```
apic1(config-tenant-13ext-epg)# contract intra-epg contr-intra
```

Adding an Intra-EPG Contract to an Application EPG Using the REST API

After you configure a contract, you can add the contract to an EPG as an intra-EPG contract. The procedure is the same for VMware VDS, OVS, and baremetal servers.

Before you begin

- You must have configured an EPG.
- You must have configured a contract with filters.

Procedure

Step 1 Configure the selectors using an XML POST request similar to the following example:

Example:

```
<?xml version="1.0" encoding="UTF-8"?>
<polUni>
  <infraInfra>
    <infraAccPortP name="Ports-1-12" status="deleted"/>

    <!-- VMM VLAN range -->
    <fvnsVlanInstP name="test" allocMode="dynamic">
      <fvnsEncapBlk name="encap" from="vlan-10" to="vlan-100"/>
    </fvnsVlanInstP>

    <!-- Static VLAN range -->
    <fvnsVlanInstP name="test" allocMode="static">
      <fvnsEncapBlk name="default" from="vlan-101" to="vlan-4095"/>
    </fvnsVlanInstP>

    <infraAttEntityP name="test">
      <infraRsDomP tDn="uni/phys-test"/>
      <infraRsDomP tDn="uni/l3dom-test"/>
      <infraRsDomP tDn="uni/vmmp-VMware/dom-test"/>
    </infraAttEntityP>

    <!-- Node profile -->
    <infraNodeP name="test">
      <infraLeafS name="test" type="range">
        <infraNodeBlk name="default" from_"=101" to_"=102"/>
      </infraLeafS>
      <infraRsAccPortP tDn="uni/infra/accportprof-test"/>
    </infraNodeP>

    <!-- Port profile -->
    <infraAccPortP name="test">
      <!-- 12 regular ports -->
      <infraHPortS name="ports1Through12" type="range">
        <infraPortBlk name="default" fromCard="1" toCard="1" fromPort="1" toPort="12"/>
        <infraRsAccBaseGrp tDn="uni/infra/funcprof/accportgrp-test"/>
      </infraHPortS>

      <!-- 2 ports in PC -->
      <infraHPortS name="portsForPc1" type="range">
        <infraPortBlk name="default" fromCard="1" toCard="1" fromPort="13" toPort="14"/>
        <infraRsAccBaseGrp tDn="uni/infra/funcprof/accbundle-testPc"/>
      </infraHPortS>
    </infraAccPortP>
  </infraInfra>
</polUni>
```

```

<!-- 2 ports in PC -->
<infraHPortS name="portsForPc2" type="range">
  <infraPortBlk name="blk1" fromCard="1" toCard="1" fromPort="15" toPort="16"/>
  <infraRsAccBaseGrp tDn="uni/infra/funcprof/accbundle-pc"/>
</infraHPortS>

<!-- 2 ports in PC for FEX -->
<infraHPortS name="portsForFex" type="range">
  <infraPortBlk name="blk1" fromCard="1" toCard="1" fromPort="17" toPort="18"/>
  <infraRsAccBaseGrp tDn="uni/infra/fexprof-default/fexbundle-test" fexId="111"/>
</infraHPortS>

<!-- 2 ports in PC for VPC -->
<infraHPortS name="portsForVpc" type="range">
  <infraPortBlk name="blk1" fromCard="1" toCard="1" fromPort="19" toPort="20"/>
  <infraRsAccBaseGrp tDn="uni/infra/funcprof/accbundle-testVpc"/>
</infraHPortS>
</infraAccPortP>

<!-- FEX profile -->
<infraFexP name="default">
  <infraFexBndlGrp name="default"/>

  <!-- 12 FEX ports -->
  <infraHPortS name="ports1Through12" type="range">
    <infraPortBlk name="default" fromCard="1" toCard="1" fromPort="1" toPort="12"/>
    <infraRsAccBaseGrp tDn="uni/infra/funcprof/accportgrp-test"/>
  </infraHPortS>

  <!-- 3 ports in FEX PC -->
  <infraHPortS name="portsForPc" type="range">
    <infraPortBlk name="blk1" fromCard="1" toCard="1" fromPort="13" toPort="16"/>
    <infraRsAccBaseGrp tDn="uni/infra/funcprof/accbundle-testPcOnFex"/>
  </infraHPortS>

  <!-- 3 ports in FEX VPC -->
  <infraHPortS name="portsForVpc" type="range">
    <infraPortBlk name="blk1" fromCard="1" toCard="1" fromPort="17" toPort="19"/>
    <infraRsAccBaseGrp tDn="uni/infra/funcprof/accbundle-testVpcOnFex"/>
  </infraHPortS>
</infraFexP>

<!-- Functional profile -->
<infraFuncP>
  <!-- Regular port group -->
  <infraAccPortGrp name="test">
    <infraRsAttEntP tDn="uni/infra/attentp-test"/>
  </infraAccPortGrp>

  <!-- PC -->
  <infraAccBndlGrp name="testPc" lagT="link">
    <infraRsLacpPol tnLacpLagPolName="testPc"/>
    <infraRsAttEntP tDn="uni/infra/attentp-test"/>
  </infraAccBndlGrp>

  <!-- VPC -->
  <infraAccBndlGrp name="testVpc" lagT="node">
    <infraRsLacpPol tnLacpLagPolName="testVpc"/>
    <infraRsAttEntP tDn="uni/infra/attentp-test"/>
  </infraAccBndlGrp>

  <!-- PC on FEX -->
  <infraAccBndlGrp name="testPcOnFex" lagT="link">

```

```

        <infraRsLacpPol tnLacpLagPolName="testPcOnFex"/>
        <infraRsAttEntP tDn="uni/infra/attentp-test"/>
    </infraAccBndlGrp>

    <!-- VPC on FEX -->
    <infraAccBndlGrp name="testVpcOnFex" lagT="node">
        <infraRsLacpPol tnLacpLagPolName="testVpcOnFex"/>
        <infraRsAttEntP tDn="uni/infra/attentp-test"/>
    </infraAccBndlGrp>
</infraFuncP>

<!-- Link aggregation policies -->
<lacpLagPol name="testPc" minLinks='1' maxLinks='10'/>
<lacpLagPol name="testVpc" minLinks='1' maxLinks='10'/>
<lacpLagPol name="testPcOnFex" minLinks='2' maxLinks='5'/>
<lacpLagPol name="testVpcOnFex" minLinks='2' maxLinks='10'/>
</infraInfra>

<fabricInst>
    <fabricProtPol name="testVpc">
        <fabricExplicitGep name="testVpc" id="101">
            <fabricNodePEp id="101"/>
            <fabricNodePEp id="102"/>
        </fabricExplicitGep>
    </fabricProtPol>
</fabricInst>

<physDomP name="test">
    <infraRsVlanNs tDn="uni/infra/vlanns-test-static"/>
</physDomP>

<l3extDomP name="test">
    <infraRsVlanNs tDn="uni/infra/vlanns-test-static"/>
</l3extDomP>
</polUni>

```

Step 2 Configure the tenant using an XML POST request similar to the following example:

Example:

```

<?xml version="1.0" encoding="UTF-8"?>
<polUni>
    <fvTenant name="Tenant-l3out">
        <vzBrCP intent="install" name="contr-intra" scope="context">
            <vzSubj consMatchT="AtleastOne" name="subj" revFltPorts="yes">
                <vzRsSubjFiltAtt action="permit" priorityOverride="default"
                    tnVzFilterName="flt-ssh" />
            </vzSubj>
        </vzBrCP>
        <vzBrCP intent="install" name="contr2" scope="context">
            <vzSubj consMatchT="AtleastOne" name="contr2-subj" revFltPorts="yes">
                <vzRsSubjFiltAtt action="permit" priorityOverride="default"
                    tnVzFilterName="flt-ftp" />
            </vzSubj>
        </vzBrCP>
        <vzBrCP intent="install" name="contr1" scope="context">
            <vzSubj consMatchT="AtleastOne" name="subj-http" revFltPorts="yes">
                <vzRsSubjFiltAtt action="deny" priorityOverride="default"
                    tnVzFilterName="flt-http" />
            </vzSubj>
        </vzBrCP>
        <l3extOut enforceRtctrl="export" mplsEnabled="no" name="l3out1">
            <l3extRsL3DomAtt tDn="uni/l3dom-test" />
            <l3extRsEctx tnFvCtxName="vrf1" />
            <l3extLNodeP name="l3out1_nodeProfile" tag="yellow-green">

```

```

<l3extRsNodeL3OutAtt rtrId="172.16.0.1" rtrIdLoopBack="yes"
  tDn="topology/pod-1/node-101" />
<l3extLIIfP name="l3out1_interfaceProfile" tag="yellow-green">
  <l3extRsPathL3OutAtt addr="192.168.15.1/24" autostate="disabled"
    encap="unknown" encapScope="local" ifInstT="l3-port" ipv6Dad="enabled"
    isMultiPodDirect="no" llAddr="::" mac="00:22:BD:F8:19:FF"
    mode="regular" mtu="inherit"
    tDn="topology/pod-1/paths-101/pathep-[eth1/10]" />
</l3extLIIfP>
</l3extLNodeP>

<!--
  Set pcEnfPref to "enforced" to enable intra-Ext-EPG isolation.
  Set pcEnfPref to "unenforced" to disable intra-Ext-EPG isolation.
-->
<l3extInstP floodOnEncap="disabled" matchT="AtleastOne"
  name="l3epg1" pcEnfPref="unenforced" prefGrMemb="exclude">
  <l3extSubnet ip="172.16.0.0/16" scope="import-security" />
  <fvRsCons tnVzBrCPName="contr2" />
  <fvRsIntraEpg tnVzBrCPName="contr-intra" />
</l3extInstP>
</l3extOut>
<fvCtx bdEnforcedEnable="no" ipDataPlaneLearning="enabled" knwMcastAct="permit"
  name="vrf1" pcEnfDir="egress" pcEnfPref="unenforced" vrfIndex="0">
  <fvRsVrfValidationPol />
  <vzAny matchT="AtleastOne" prefGrMemb="disabled" />
</fvCtx>
<fvBD OptimizeWanBandwidth="no" arpFlood="yes" epClear="no" hostBasedRouting="no"
  intersiteBumTrafficAllow="no" intersiteL2Stretch="no" ipLearning="yes"
  ipv6McastAllow="no" limitIpLearnToSubnets="yes" llAddr="::"
  mac="00:22:BD:F8:19:FF" mcastAllow="no" multiDstPktAct="bd-flood" name="bd-web"
  type="regular" unicastRoute="yes" unkMacUcastAct="proxy" unkMcastAct="flood"
  v6unkMcastAct="flood" vmac="not-applicable">
  <fvSubnet ip="192.168.1.254/24" ipDPLearning="enabled" preferred="no"
    scope="private" virtual="no" />
  <fvRsCtx tnFvCtxName="vrf1" />
  <fvRsBdToEpRet resolveAct="resolve" />
</fvBD>
<fvBD OptimizeWanBandwidth="no" arpFlood="yes" epClear="no" hostBasedRouting="no"
  intersiteBumTrafficAllow="no" intersiteL2Stretch="no" ipLearning="yes"
  ipv6McastAllow="no" limitIpLearnToSubnets="yes" llAddr="::"
  mac="00:22:BD:F8:19:FF" mcastAllow="no" multiDstPktAct="bd-flood" name="bd-app"
  type="regular" unicastRoute="yes" unkMacUcastAct="proxy" unkMcastAct="flood"
  v6unkMcastAct="flood" vmac="not-applicable">
  <fvSubnet ip="192.168.2.254/24" ipDPLearning="enabled" preferred="no"
    scope="private" virtual="no" />
  <fvRsCtx tnFvCtxName="vrf1" />
  <fvRsBdToEpRet resolveAct="resolve" />
</fvBD>
<vzFilter name="flt-ftp">
  <vzEntry applyToFrag="no" arpOpc="unspecified" dFromPort="ftpData"
    dToPort="ftpData" etherT="ipv4" icmpv4T="unspecified" icmpv6T="unspecified"
    matchDscp="unspecified" name="ftp" prot="tcp" sFromPort="unspecified"
    sToPort="unspecified" stateful="no" />
</vzFilter>
<vzFilter name="flt-ssh">
  <vzEntry applyToFrag="no" arpOpc="unspecified" dFromPort="ssh" dToPort="ssh"
    etherT="ipv4" icmpv4T="unspecified" icmpv6T="unspecified"
    matchDscp="unspecified" name="ssh" prot="tcp" sFromPort="unspecified"
    sToPort="unspecified" stateful="no" />
</vzFilter>
<vzFilter name="flt-http">
  <vzEntry applyToFrag="no" arpOpc="unspecified" dFromPort="http" dToPort="http"
    etherT="ipv4" icmpv4T="unspecified" icmpv6T="unspecified"

```

```

        matchDscp="unspecified" name="flt1" prot="tcp" sFromPort="unspecified"
        sToPort="unspecified" stateful="no" />
    </vzFilter>
    <fvAp name="ap-appl">
        <fvAEPg floodOnEncap="disabled" hasMcastSource="no" isAttrBasedEPg="no"
        matchT="AtleastOne" name="epg-app" pcEnfPref="unenforced"
        prefGrMemb="exclude" shutdown="no">
            <fvRsProv intent="install" matchT="AtleastOne" tnVzBrCPName="contr2" />
            <fvRsProv intent="install" matchT="AtleastOne" tnVzBrCPName="contr1" />
            <fvRsPathAtt encap="vlan-103" instrImedcy="immediate" mode="native"
            primaryEncap="unknown" tDn="topology/pod-1/paths-101/pathep-[eth1/3]" />
            <fvRsDomAtt bindingType="none" classPref="encap" encap="unknown"
            encapMode="auto" epgCos="Cos0" epgCosPref="disabled" instrImedcy="lazy"
            netflowDir="both" netflowPref="disabled" numPorts="0" portAllocation="none"
            primaryEncap="unknown" primaryEncapInner="unknown" resImedcy="immediate"
            secondaryEncapInner="unknown" switchingMode="native" tDn="uni/phys-test"
            untagged="no" vnetOnly="no" />
            <fvRsBd tnFvBDName="bd-app" />
        </fvAEPg>
        <fvAEPg floodOnEncap="disabled" hasMcastSource="no"
        isAttrBasedEPg="no" matchT="AtleastOne" name="epg-web" pcEnfPref="unenforced"
        prefGrMemb="exclude" shutdown="no">
            <fvRsPathAtt encap="vlan-104" instrImedcy="immediate" mode="native"
            primaryEncap="unknown" tDn="topology/pod-1/paths-101/pathep-[eth1/4]" />
            <fvRsDomAtt bindingType="none" classPref="encap" encap="unknown"
            encapMode="auto" epgCos="Cos0" epgCosPref="disabled" instrImedcy="lazy"
            netflowDir="both" netflowPref="disabled" numPorts="0" portAllocation="none"
            primaryEncap="unknown" primaryEncapInner="unknown" resImedcy="immediate"
            secondaryEncapInner="unknown" switchingMode="native" tDn="uni/phys-test"
            untagged="no" vnetOnly="no" />
            <fvRsCons intent="install" tnVzBrCPName="contr1" />
            <fvRsBd tnFvBDName="bd-web" />
        </fvAEPg>
    </fvAp>
</fvTenant>
</polUni>

```

EPG Contract Inheritance

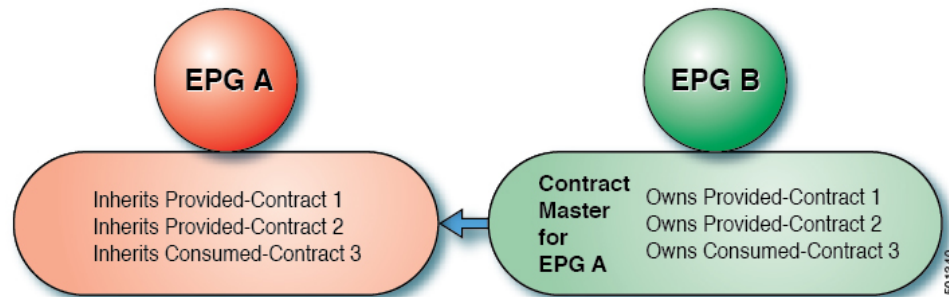
About Contract Inheritance

To streamline associating contracts to new EPGs, you can now enable an EPG to inherit all the (provided and consumed) contracts associated directly to another EPG in the same tenant. Contract inheritance can be configured for application, microsegmented, L2Out, and L3Out EPGs.

With Release 3.x, you can also configure contract inheritance for Inter-EPG contracts, both provided and consumed. Inter-EPG contracts are supported on Cisco Nexus 9000 Series switches with EX or FX at the end of their model name or later models.

You can enable an EPG to inherit all the contracts associated directly to another EPG, using the APIC GUI, NX-OS style CLI, and the REST API.

Figure 9: Contract Inheritance



In the diagram above, EPG A is configured to inherit Provided-Contract 1 and 2 and Consumed-Contract 3 from EPG B (contract master for EPG A).

Use the following guidelines when configuring contract inheritance:

- Contract inheritance can be configured for application, microsegmented (uSeg), external L2Out EPGs, and external L3Out EPGs. The relationships must be between EPGs of the same type.
- Both provided and consumed contracts are inherited from the contract master when the relationship is established.
- Contract masters and the EPGs inheriting contracts must be within the same tenant.
- Changes to the masters' contracts are propagated to all the inheritors. If a new contract is added to the master, it is also added to the inheritors.
- An EPG can inherit contracts from multiple contract masters.
- Contract inheritance is only supported to a single level (cannot be chained) and a contract master cannot inherit contracts.
- Labels with contract inheritance is supported. When EPG A inherits a contract from EPG B, if different subject labels are configured under EPG A and EPG B, APIC uses the label configured under EPG B for the contract inherited from EPG B. APIC uses the label configured under EPG A for the contract where EPG A is directly involved.
- Whether an EPG is directly associated to a contract or inherits a contract, it consumes entries in TCAM. So contract scale guidelines still apply. For more information, see the *Verified Scalability Guide* for your release.
- vzAny security contracts and taboo contracts are not supported.
- Beginning in Cisco APIC releases 5.0(1) and 4.2(6), contract inheritance with a service graph is supported if the contract and EPGs are in the same tenant.

For information about configuring Contract Inheritance and viewing inherited and standalone contracts, see *Cisco APIC Basic Configuration Guide*.

Configuring EPG Contract Inheritance Using the GUI

Configuring Application EPG Contract Inheritance Using the GUI

To configure contract inheritance for an application EPG, in the APIC Basic or Advanced mode GUI, use the following steps.

Before you begin

Configure the tenant and application profile to be used by the EPGs.

Optional. Configure the bridge domain to be used by the EPG that will inherit contracts.

Configure at least one application EPG, to serve as the **EPG Contract Master**.

Configure the contracts to be shared, and associate them to the contract master.

Procedure

-
- Step 1** Navigate to **Tenants > *tenant-name* > Application Profiles**, and expand ***AP-name***
 - Step 2** Right-click **Application EPGs** and select **Create Application EPG**.
 - Step 3** Type the name of the EPG that will inherit contracts from the **EPG Contract Master**.
 - Step 4** On the **Bridge Domain** field, select the common/default bridge domain or a previously created bridge domain, or create a bridge domain for this EPG.
 - Step 5** On the **EPG Contract Master** field, click the + symbol, select the previously configured Application Profile and EPG, and click **Update**.
 - Step 6** Click **Finish**.
 - Step 7** To view information about the EPG, including the contract master, navigate to **Tenants > *tenant-name* > Application Profiles > *AP-name* > Application EPGs > *EPG-name***. To view the **EPG Contract Master**, click **General**.
 - Step 8** To view information about the inherited contracts, expand ***EPG-name*** and click **Contracts**.
-

Configuring uSeg EPG Contract Inheritance Using the GUI

To configure contract inheritance for a uSeg EPG, in the APIC Basic or Advanced mode GUI, use the following steps.

Before you begin

Configure the tenant and application profile to be used by the EPGs.

Optional. Configure the bridge domain to be used by the EPG that will inherit contracts.

Configure the uSeg EPG, to serve as the **EPG Contract Master**.

Configure the contracts to be shared, and associate them to the contract master.

Procedure

-
- Step 1** Navigate to **Tenants** > *tenant-name* > **Application Profiles**, expand *AP-name*.
 - Step 2** Right-click **uSeg EPGs** and select **Create uSeg EPG**.
 - Step 3** Type the name of the EPG that will inherit contracts from the contract master.
 - Step 4** On the **Bridge Domain** field, select the common/default bridge domain or a previously created bridge domain, or create a bridge domain for this EPG.
 - Step 5** Click *uSeg-EPG-name*. In the **EPG Contract Master** field, click the + symbol, select the Application Profile and EPG (to serve as contract master), and click **Update**.
 - Step 6** Click **Finish**.
 - Step 7** To view information about the contracts, navigate to **Tenants** > *tenant-name* > **Application Profiles** > *AP-name* > **uSeg EPGs** > , expand the *EPG-name* and click **Contracts**.
-

Configuring L2Out EPG Contract Inheritance Using the GUI

To configure contract inheritance for an external L2Out EPG, in the Cisco Application Policy Infrastructure Controller (APIC) GUI, perform the following steps.

Before you begin

Configure the tenant and application profile to be used by the EPGs.

Configure a Layer 2 Outside (L2Out) and the external L2Out EPG (L2extInstP) that will serve as the **L2Out Contract Master**.

Configure the contracts to be shared, and associate them to the contract master.

Procedure

-
- Step 1** Navigate to **Tenants** > *tenant-name* > **Networking** > **L2Outs**.
 - Step 2** Expand the *L2Out-name*.
 - Step 3** Right-click **External EPGs** and choose **Create External EPG**.
 - Step 4** Type the name of the external network and optionally add other attributes.
 - Step 5** Click **Submit**.
 - Step 6** Expand **External EPGs**.
 - Step 7** Click the *external-epg-name*.
 - Step 8** In the **External EPG** panel, click the + symbol on the **L2Out Contract Masters** field.
 - Step 9** Select the L2Out and the L2Out contract master for this external L2Out EPG.
 - Step 10** Click **Update**.
 - Step 11** To view the contracts inherited by this external L2Out EPG, click on the external EPG name and click **Contracts** > **Inherited Contracts**.
-

Configuring External L3Out EPG Contract Inheritance Using the GUI

To configure contract inheritance for an external L3Out EPG, in the Cisco Application Policy Infrastructure Controller (APIC) GUI, use the following steps.

Before you begin

Configure the tenant and application profile to be used by the EPGs.

Configure an external routed network (L3Out) and the external L3Out EPG (L3extInstP) that will serve as the **L3Out Contract Master**.

Configure the contracts to be shared, and associate them to the contract master.

Procedure

-
- Step 1** To configure contract inheritance for an external L3Out EPG, navigate to **Tenants > *tenant-name* > Networking > L3Outs**.
 - Step 2** Expand the ***L3Out-name*** leading to the external L3Out EPG.
 - Step 3** Right-click **External EPGs** and select **Create External EPG**.
 - Step 4** Type the name of the external EPG and optionally add subnets and other attributes.
 - Step 5** Click **Submit**.
 - Step 6** Expand **Networks**.
 - Step 7** Click the ***network-name***.
 - Step 8** In the **External EPG** panel, click the + symbol on the **L3Out Contract Masters** field.
 - Step 9** Choose the L3Out and external EPG to serve as L3Out contract master for this external L3Out EPG.
 - Step 10** Click **Update**.
 - Step 11** To view the contracts inherited by this external L3Out EPG, click on the external EPG name and click **Contracts > Inherited Contracts**.
-

Contract Preferred Groups

About Contract Preferred Groups

There are two types of policy enforcements available for EPGs in a VRF with a contract preferred group configured:

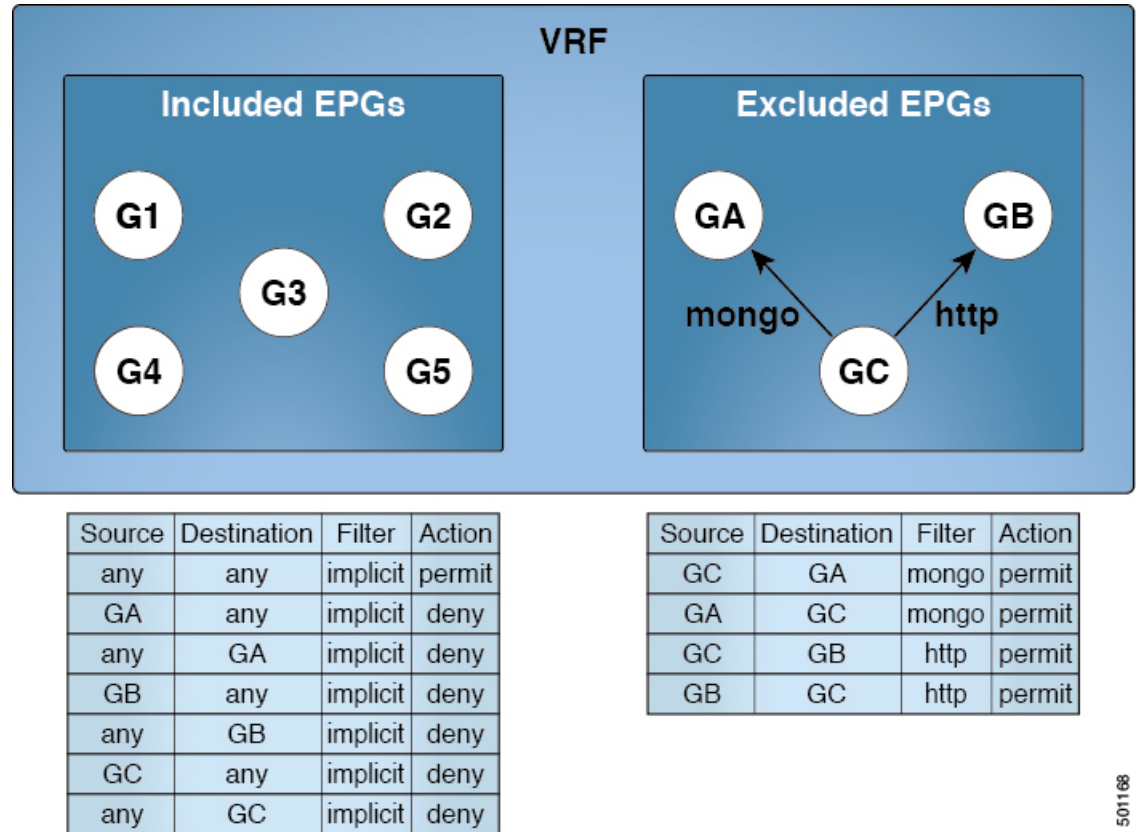
- **Included EPGs:** EPGs can freely communicate with each other without contracts, if they have membership in a contract preferred group. This is based on the source-any-destination-any-permit default rule.
- **Excluded EPGs:** EPGs that are not members of preferred groups require contracts to communicate with each other. Otherwise, the default source-any-destination-any-deny rule applies.

The contract preferred group feature enables greater control of communication between EPGs in a VRF. If most of the EPGs in the VRF should have open communication, but a few should only have limited

communication with the other EPGs, you can configure a combination of a contract preferred group and contracts with filters to control inter-EPG communication precisely.

EPGs that are excluded from the preferred group can only communicate with other EPGs if there is a contract in place to override the source-any-destination-any-deny default rule.

Figure 10: Contract Preferred Group Overview



501168

Limitations

The following limitations apply to contract preferred groups:

- In topologies where an L3Out and application EPG are configured in a Contract Preferred Group, and the EPG is deployed only on a VPC, you may find that only one leaf switch in the VPC has the prefix entry for the L3Out. In this situation, the other leaf switch in the VPC does not have the entry, and therefore drops the traffic.

To workaround this issue, you can do one of the following:

- Disable and reenab the contract group in the VRF
- Delete and recreate the prefix entries for the L3Out EPG
- Also, where the provider or consumer EPG in a service graph contract is included in a contract group, the shadow EPG can not be excluded from the contract group. The shadow EPG will be permitted in the contract group, but it does not trigger contract group policy deployment on the node where the shadow

EPG is deployed. To download the contract group policy to the node, you deploy a dummy EPG within the contract group .

Due to CSCvm63145, an EPG in a Contract Preferred Group can consume a shared service contract, but cannot be a provider for a shared service contract with an L3Out EPG as consumer.

Guidelines for Contract Preferred Groups

When configuring contract preferred groups, refer to the following guidelines:

- If the (s, g) entry is installed on a border leaf switch, you might see drops in unicast traffic that comes from the fabric to this source outside the fabric when the following conditions are met:
 - Preferred group is used on the L3Out EPG
 - Unicast routing table for the source is using the default route 0.0.0.0/0

This behavior is expected.

- Contract Preferred Group-included EPGs are not supported with a 0/0 prefix in external EPG (InstP). If, for the external EPG (InstP) to Tenant EPG, a 0/0 prefix is required with the use of Contract Preferred Group, then 0/0 can be split to 0/1 and 128/1.
- Contract Preferred Group-EPGs are not supported with the GOLF feature. Communication between an application EPG and the L3Out EPG for GOLF must be governed by explicit contracts.

Configuring Contract Preferred Groups Using the GUI

Before you begin

Create the tenants and VRF, and EPGs that will consume the contract preferred group.

Procedure

-
- | | |
|----------------|---|
| Step 1 | On the menu bar, click Tenants > <i>tenant_name</i> . |
| Step 2 | In the Navigation pane, expand the tenant, Networking , and VRFs . |
| Step 3 | Click the VRF name for which you are configuring the contract preferred group. |
| Step 4 | In the Preferred Group Member field, click Enabled . |
| Step 5 | Click Submit . |
| Step 6 | In the Navigation pane, expand Application Profiles and create or expand an application profile for the tenant VRF. |
| Step 7 | Expand Application EPGs and click the EPG that will consume the contract preferred group. |
| Step 8 | Select the Policy and General tab. |
| Step 9 | In the Preferred Group Member field, click Include . |
| Step 10 | Click Submit . |
-

What to do next

Enable membership in the preferred group for other EPGs that should have unlimited communication with this EPG. You can also configure appropriate contracts to control communication between the EPGs in the preferred group and other EPGs that may not be members.



Note If you want to support preferred group members through L4-L7 service graphs, you must create a L4-L7 service EPG policy. For more information regarding creating an L4-L7 Service EPG Policy, see [Creating an L4-L7 Service EPG Policy Using the GUI, on page 231](#).

Creating an L4-L7 Service EPG Policy Using the GUI

This task creates a policy that defines if EPGs are to be included in, or excluded from, a preferred group. Preferred groups membership allows endpoints to communicate with each other without requiring a contract. After the policy is created, it can be selected during the application of a service graph template to the EPGs.

Before you begin

You must have configured a tenant.

Procedure

-
- Step 1** On the Menu bar, choose **Tenants** > *tenant_name*.
- Step 2** In the Navigation pane, choose **Policies** > **Protocol** > **L4-L7 Service EPG Policy**.
- Step 3** In the Navigation pane, right-click **L4-L7 Service EPG Policy** and choose **Create L4-L7 Service EPG Policy**.
- The Create L4-L7 Service EPG Policy dialog box appears.
- Step 4** Enter a unique name for the policy in the **Name** field.
- Step 5** Optional. Enter a description of the policy in the **Description** field.
- Step 6** Choose whether to exclude or include EPGs as preferred members in the **Preferred Group Member** field.
- Step 7** Click **Submit**.
- The newly created policy appears in the L4-L7 Service EPG Policy Work pane list. To edit a policy in the Work pane, double-click the list line containing the policy.
-

What to do next

The new L4-L7 service EPG policy can now be selected in a service graph template when applying the graph to EPGs. Refer to [Applying a Service Graph Template to Endpoint Groups Using the GUI](#) in the Using the GUI chapter of the *Cisco APIC Layer 4 to Layer 7 Services Deployment Guide*.

Contracts with Permit and Deny Rules

About Contracts with Permit and Deny Rules

Starting with the Cisco Application Policy Infrastructure Controller (Cisco APIC) release 3.2, You can configure contracts with both permit and deny actions, instead of just permit. You can configure the deny action with different priorities: default, highest, medium and lowest.

Rule conflicts are resolved as follows:

- The implicit deny has the lowest priority of all rules.
- Contracts between vzAny have higher priority than the implicit deny.
- Contracts between specific EPG pairs win over contracts with vzAny, because EPG-to-EPG contract rules have higher priority than vzAny-to-vzAny rules.
- Deny rules with the default priority for a contract between a specific EPG pair have the same level of priority as the permit rules for that EPG pair. When traffic matches both a permit and a deny rule with the same priority, the deny rule wins.
- Deny rules with the default priority for a contract between vzAny has the same level of priority as the permit rules for the vzAny pair. When traffic matches both a permit and a deny rule with the same priority, the deny rule wins.
- Deny rules with the highest priority are handled at the same level as EPG-to-EPG contracts.
- Deny rules with medium priority are handled at the same level as vzAny-to-EPG contracts.
- Deny rules with the lowest priority are handled at the same level as vzAny-to-vzAny contracts.
- If the deny priority is lowered in a contract between EPGs, a permit rule match between EPGs would win over deny.

Enabling ACL Contract Permit and Deny Logging Using the GUI

The following steps show how to enable contract permit and deny logging using the GUI:



Note The tenant that contains the permit logging is the tenant that contains the VRF that the EPG is associated to. This will not necessarily be the same tenant as the EPG or its associated contracts.

Procedure

-
- Step 1** On the menu bar, choose **Tenants** > <tenant name>.
 - Step 2** In the **Navigation** pane, expand **Contracts**, right-click **Standard**, and choose **Create Contract**.
 - Step 3** In the **Create Contract** dialog box, perform the following actions:

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

- Step 4** In the Create Contract Subject dialog box, perform the following actions:
- Step 5** Enter the name of the subject and an optional description.
- Step 6** Optional. From the drop-down list for the target DSCP, select the DSCP to be applied to the subject.
- Step 7** Leave **Apply Both Directions** checked, unless you want the contract to only be applied from the consumer to the provider, instead of in both directions.
- Step 8** Leave **Reverse Filter Ports** checked if you unchecked **Apply Both Directions** to swap the Layer 4 source and destination ports so that the rule is applied from the provider to the consumer.
- Step 9** Click the + icon to expand **Filters**.
- Step 10** In the **Name** drop-down list, choose an option; for example, click **arp**, **default**, **est**, or **icmp**, or choose a previously configured filter.
- Step 11** In the **Directives** drop-down list, click **log**.
- Step 12** (Optional) Change the Action to be taken with this subject to **Deny** (or leave the action to the default, **Permit**).
With Directive: log enabled, if the action for this subject is **Permit**, ACL permit logs track the flows and packets that are controlled by the subject and contract. If the action for this subject is **Deny**, ACL deny logs track the flows and packets.
- Step 13** (Optional) Set the priority for the subject.
- Step 14** Click **Update**.
- Step 15** Click **OK**.
- Step 16** Click **Submit**.
Logging is enabled for this contract.
-

Enabling ACL Contract Permit Logging Using the NX-OS CLI

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

Procedure

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

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

Example:

For example:

```
apic1# configure
apic1(config)# tenant BDMoDel
```

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

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

Enabling ACL Contract Permit Logging Using the REST API

The following example shows you how to enable permit and deny logging using the REST API. This example configures ACL permit and deny logging for a contract with subjects that have Permit and Deny actions configured.

Procedure

For this configuration, send a post with XML similar to the following example:

Example:

```
<vzBrCP dn="uni/tn-Tenant64/brc-C64" name="C64" scope="context">
  <vzSubj consMatchT="AtleastOne" name="HTTPSsbj" provMatchT="AtleastOne" revFltPorts="yes"
    rn="subj-HTTPSsbj">
    <vzRsSubjFiltAtt action="permit" directives="log" forceResolve="yes"
      priorityOverride="default"
      rn="rbsubjFiltAtt-PerHTTPS" tDn="uni/tn-Tenant64/flt-PerHTTPS" tRn="flt-PerHTTPS"
      tnVzFilterName="PerHTTPS"/>
    </vzSubj>
    <vzSubj consMatchT="AtleastOne" name="httpSbj" provMatchT="AtleastOne" revFltPorts="yes"
      rn="subj-httpSbj">
      <vzRsSubjFiltAtt action="deny" directives="log" forceResolve="yes"
        priorityOverride="default"
        rn="rbsubjFiltAtt-httpFilter" tDn="uni/tn-Tenant64/flt-httpFilter" tRn="flt-httpFilter"
        tnVzFilterName="httpFilter"/>
      </vzSubj>
      <vzSubj consMatchT="AtleastOne" name="subj64" provMatchT="AtleastOne" revFltPorts="yes"
        rn="subj-subj64">
        <vzRsSubjFiltAtt action="permit" directives="log" forceResolve="yes"
          priorityOverride="default"
          rn="rbsubjFiltAtt-icmp" tDn="uni/tn-common/flt-icmp" tRn="flt-icmp" tnVzFilterName="icmp"/>
        </vzSubj>
      </vzBrCP>
```

Enabling Taboo Contract Deny Logging Using the GUI

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

Procedure

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

- Step 2** In the **Navigation** pane, expand **Contracts**.
- Step 3** Right-click **Taboos** and choose **Create Taboo Contract**.
- Step 4** In the Create Taboo Contract dialog box, perform the following actions to specify the Taboo contract:
- In the **Name** field, type the name for the contract.
 - Optional. In the **Description** field, type a description of the Taboo contract.
 - Click the + icon to expand **Subjects**.
- Step 5** In the **Create Taboo Contract Subject** dialog box, perform the following actions:
- In the Specify Identity of Subject area, type a name and optional description.
 - Click the + icon to expand **Filters**.
 - From the **Name** drop-down list, choose one of the default values, such as `<tenant_name>/arp`, `<tenant_name>/default`, `<tenant_name>/est`, `<tenant_name>/icmp`, choose a previously created filter, or **Create Filter**.
- Note**
If you chose **Create Filter**, in the Specify Filter Identity Area, perform the following actions to specify criteria for the ACL Deny rule:
- Type a name and optional description.
 - Expand **Entries**, type a name for the rule, and choose the criteria to define the traffic you want to deny.
 - In the Directives drop-down list, choose **log**.
 - Click **Update**.
 - Click **OK**.
- Step 6** Click **Submit**.
Logging is enabled for this Taboo contract.

Enabling Taboo Contract Deny Logging Using the NX-OS CLI

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

Procedure

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

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

Example:

For example:

```
apic1# configure
apic1(config)# tenant BDMoDel
```

```
apicl(config-tenant)# contract dropFTP type deny
apicl(config-tenant-contract)# subject dropftp
apicl(config-tenant-contract-subj)# access-group ftp both log
```

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

Enabling Taboo Contract Deny Logging Using the REST API

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

Procedure

To configure taboo contract deny logging, send a post with XML similar to the following example.

Example:

```
<vzTaboo dn="uni/tn-Tenant64/taboo-TCtrctPrefix" name="TCtrctPrefix" scope="context">
  <vzTSubj name="PrefSubj" rn="tsubj-PrefSubj">
    <vzRsDenyRule directives="log" forceResolve="yes" rn="rsdenyRule-default"
    tCl="vzFilter"
    tDn="uni/tn-common/flt-default" tRn="flt-default"/>
  </vzTSubj>
</vzTaboo>
```

Viewing ACL Permit and Deny Logs Using the GUI

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

Procedure

- Step 1** On the menu bar, choose **Tenants > <tenant name>**.
- Step 2** In the **Navigation** pane, click on **Tenant <tenant name>**.
- Step 3** In the **Tenants <tenant name> Work** pane, click the **Operational** tab.
- Step 4** Under the **Operational** tab, click the **Flows** tab.
Under the **Flows** tab, click one of the tabs to view log data for Layer 2 permit logs (**L2 Permit**) Layer 3 permit logs (**L3 Permit**, Layer 2 deny logs (**L2 Drop**), or Layer 3 deny logs (**L3 Drop**). On each tab, you can view ACL logging data, if traffic is flowing. The data points differ according to the log type and ACL rule; for example, the following data points are included for **L3 Permit** and **L3 Deny** logs:
- VRF
 - Alias
 - Source IP address

- Destination IP address
- Protocol
- Source port
- Destination port
- Source MAC address
- Destination MAC address
- Node
- Source interface
- VRF Encap
- Source EPG
- Destination EPG
- Source PC Tag
- Destination PC Tag

Note

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

Viewing ACL Permit and Deny Logs Using the REST API

The following example shows how to view Layer 2 deny log data for traffic flows, using the REST API. You can send queries using the following MOs:

- aclogDropL2Flow
- aclogPermitL2Flow
- aclogDropL3Flow
- aclogPermitL3Flow
- aclogDropL2Pkt
- aclogPermitL2Pkt
- aclogDropL3Pkt
- aclogPermitL3Pkt

Before you begin

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

Procedure

To view Layer 3 drop log data, send the following query using the REST API:

```
GET https://apic-ip-address/api/class/aclogDropL3Flow
```

Example:

The following example shows sample output:

```
<?xml version="1.0" encoding="UTF-8"?>
<imdata totalCount="2">
  <aclogPermitL3Flow childAction=""
dn="topology/pod-1/node-101/ndbgs/aclog/tn-common/ctx-inb
/permitl3flow-spctag-333-dpctag-444-sepname-unknown-depname-unknown-sip-[100:c000:a00:700:b00:0:f00:0]
-dip-[19.0.2.10]-proto-udp-sport-17459-dport-8721-smac-00:00:15:00:00:28-dmac-00:00:12:00:00:25-sintf-
[port-channel5]-vrfencap-VXLAN: 2097153" dstEpgName="unknown" dstIp="19.0.2.10"
dstMacAddr="00:00:12:00:00:25"
dstPcTag="444" dstPort="8721" lcOwn="local" modTs="never" monPolDn="" protocol="udp"
srcEpgName="unknown"
srcIntf="port-channel5" srcIp="100:c000:a00:700:b00:0:f00:0"
srcMacAddr="00:00:15:00:00:28" srcPcTag="333"
srcPort="17459" status="" vrfEncap="VXLAN: 2097153"/>
  <aclogPermitL3Flow childAction=""
dn="topology/pod-1/node-102/ndbgs/aclog/tn-common/ctx-inb
/permitl3flow-spctag-333-dpctag-444-sepname-unknown-depname-unknown-sip-[100:c000:a00:700:b00:0:f00:0]-dip-
[19.0.2.10]-proto-udp-sport-17459-dport-8721-smac-00:00:15:00:00:28-dmac-00:00:12:00:00:25-sintf-
[port-channel5]-vrfencap-VXLAN: 2097153" dstEpgName="unknown" dstIp="19.0.2.10"
dstMacAddr="00:00:12:00:00:25"
dstPcTag="444" dstPort="8721" lcOwn="local" modTs="never" monPolDn="" protocol="udp"
srcEpgName="unknown"
srcIntf="port-channel5" srcIp="100:c000:a00:700:b00:0:f00:0"
srcMacAddr="00:00:15:00:00:28" srcPcTag="333"
srcPort="17459" status="" vrfEncap="VXLAN: 2097153"/>
</imdata>
```

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

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

The syntax for the Layer 3 command is **show aclog {permit | deny} l3 {pkt | flow} tenant <tenant_name> vrf <vrf_name> srcip <source_ip> dstip <destination_ip> srcport <source_port> dstport <destination_port> protocol <protocol> srcintf <source_interface> start-time <startTime> end-time <endTime> detail**

The syntax for the Layer 2 command is **show aclog {permit | deny} l2 {flow | pkt} tenant <tenant_name> vrf <VRF_name> srcintf <source_interface> vlan <VLAN_number> detail**



Note The full syntax of the **show acllog** command is only available on Generation 2 Cisco Nexus 9000 series switches (with names that end in EX or FX or later, such as N9K-C93180LC-EX) and Cisco APIC Release 3.2 or later. With Generation 1 switches (with names that do not end in EX or FX) or Cisco APIC releases before 3.2, the available syntax is as above.

In Cisco APIC 3.2 and later, additional keywords are added to both versions of the command, with the **detail** keyword: **[dstEpgName <destination_EPG_name>| dstmac <destination_MAC_address> | dstpctag <destination_PCTag> | srcEpgName <source_EPG_name> | srcmac <source_MAC_address> | srcpctag <source_PCTag>]**

Procedure

Step 1 The following example shows how to use the **show acllog drop l3 flow tenant common vrf default detail** command to display detailed information about Layer 3 deny logs for the common tenant:

Example:

```
apic1# show acllog deny l3 flow tenant common vrf default detail
SrcPcTag   : 49153
DstPcTag   : 32773
SrcEPG     : uni/tn-TSW_Tenant0/ap-tsw0AP0/epg-tsw0ctx0BD0epg6
DstEPG     : uni/tn-TSW_Tenant0/ap-tsw0AP0/epg-tsw0ctx0BD0epg5
SrcIp      : 16.0.2.10
DstIp      : 19.0.2.10
Protocol   : udp
SrcPort    : 17459
DstPort    : 8721
SrcMAC     : 00:00:15:00:00:28
DstMAC     : 00:00:12:00:00:25
Node       : 101
SrcIntf    : port-channel15
VrfEncap   : VXLAN: 2097153
```

This example shows the output on Generation 2 switches, with Cisco APIC Release 3.2 or later.

Step 2 The following example shows how to use the **show acllog deny l2 flow tenant common vrf tsw0connctx0 detail** command to display detailed information about Layer 3 deny logs for the common tenant:

Example:

```
apic1# show acllog deny l2 flow tenant common vrf tsw0connctx0 detail
SrcPcTag  DstPcTag  SrcEPG          DstEPG          SrcMAC          DstMAC          Node
SrcIntf   vlan
-----
-----
-----
32773     49153     uni/tn-TSW      uni/tn-TSW      00:00:11:00:00:11  11:00:32:00:00:33  101
port-     2
channel8  _Tenant0/ap-  _Tenant0/ap-
          tsw0AP0/epg-  tsw0AP0/epg-
          tsw0ctx0BD0epg5  tsw0ctx0BD0epg6
```

This example shows the output on Generation 2 switches, with Cisco APIC Release 3.2 or later.

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

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

This example shows the output on Generation 1 switches, or with Cisco APIC releases before 3.2.

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

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

This example shows the output on Generation 1 switches, or with Cisco APIC releases before 3.2.



APPENDIX **A**

Configuring the Cisco APIC Using the CLI

- [Cluster Management Guidelines](#), on page 241
- [Replacing a Cisco APIC in a Cluster Using the CLI](#), on page 242
- [Reducing the APIC Cluster Size](#), on page 244
- [Contracting the Cisco APIC Cluster](#), on page 244
- [Switching Over Active APIC with Standby APIC Using CLI](#), on page 245
- [Verifying Cold Standby Status Using the CLI](#), on page 245
- [Registering an Unregistered Switch Using the CLI](#), on page 246
- [Adding a Switch Before Discovery Using the CLI](#), on page 246
- [Removing a Switch to Maintenance Mode Using the CLI](#), on page 247
- [Inserting a Switch to Operation Mode Using the CLI](#), on page 247
- [Configuring a Remote Location Using the NX-OS Style CLI](#), on page 247
- [Finding Your Switch Inventory Using the NX-OS CLI](#), on page 248
- [Verifying the Cisco APIC Cluster Using the CLI](#), on page 251

Cluster Management Guidelines

The Cisco Application Policy Infrastructure Controller (APIC) cluster comprises multiple Cisco APICs that provide operators a unified real time monitoring, diagnostic, and configuration management capability for the Cisco Application Centric Infrastructure (ACI) fabric. To assure optimal system performance, use the following guidelines when making changes to the Cisco APIC cluster:

- Prior to initiating a change to the cluster, always verify its health. When performing planned changes to the cluster, all controllers in the cluster should be healthy. If one or more of the Cisco APICs' health status in the cluster is not "fully fit," remedy that situation before proceeding. Also, assure that cluster controllers added to the Cisco APIC are running the same version of firmware as the other controllers in the Cisco APIC cluster.
- We recommend that you have at least 3 active Cisco APICs in a cluster, along with additional standby Cisco APICs. Cisco APIC clusters can have from 3 to 7 active Cisco APICs. Refer to the [Verified Scalability Guide](#) to determine how many active Cisco APICs are required for your deployment.
- Disregard cluster information from Cisco APICs that are not currently in the cluster; they do not provide accurate cluster information.
- Cluster slots contain a Cisco APIC `ChassisID`. Once you configure a slot, it remains unavailable until you decommission the Cisco APIC with the assigned `ChassisID`.

- If a Cisco APIC firmware upgrade is in progress, wait for it to complete and the cluster to be fully fit before proceeding with any other changes to the cluster.
- When moving a Cisco APIC, first ensure that you have a healthy cluster. After verifying the health of the Cisco APIC cluster, choose the Cisco APIC that you intend to shut down. After the Cisco APIC has shut down, move the Cisco APIC, re-connect it, and then turn it back on. From the GUI, verify that the all controllers in the cluster return to a fully fit state.



Note Only move one Cisco APIC at a time.

- When moving a Cisco APIC that is connected to a set of leaf switches to another set of leaf switches or when moving a Cisco APIC to different port within the same leaf switch, first ensure that you have a healthy cluster. After verifying the health of the Cisco APIC cluster, choose the Cisco APIC that you intend to move and decommission it from the cluster. After the Cisco APIC is decommissioned, move the Cisco APIC and then commission it.
- Before configuring the Cisco APIC cluster, ensure that all of the Cisco APICs are running the same firmware version. Initial clustering of Cisco APICs running differing versions is an unsupported operation and may cause problems within the cluster.
- Unlike other objects, log record objects are stored only in one shard of a database on one of the Cisco APICs. These objects get lost forever if you decommission or replace that Cisco APIC.
- When you decommission a Cisco APIC, the Cisco APIC loses all fault, event, and audit log history that was stored in it. If you replace all Cisco APICs, you lose all log history. Before you migrate a Cisco APIC, we recommend that you manually backup the log history.

Replacing a Cisco APIC in a Cluster Using the CLI



-
- Note**
- For more information about managing clusters, see [Cluster Management Guidelines, on page 241](#).
 - When you replace an Cisco APIC, the password will always be synced from the cluster. When replacing APIC 1, you will be asked for a password but it will be ignored in favor of the existing password in the cluster. When replacing Cisco APIC 2 or 3, you will not be asked for a password.
-

Before you begin

Before replacing a Cisco Application Policy Infrastructure Controller (APIC), ensure that the replacement Cisco APIC is running the same firmware version as the Cisco APIC to be replaced. If the versions are not the same, you must update the firmware of the replacement Cisco APIC before you begin. Initial clustering of Cisco APICs running differing versions is an unsupported operation and may cause problems within the cluster.

Procedure

-
- Step 1** Identify the Cisco APIC that you want to replace.
- Step 2** Note the configuration details of the Cisco APIC to be replaced by using the **acdiag avread** command.
- Step 3** Decommission the Cisco APIC using the **decommission controller** *controller-id* command in config mode.
- Decommissioning the Cisco APIC removes the mapping between the APIC ID and Chassis ID. The new Cisco APIC typically has a different APIC ID, so you must remove this mapping in order to add a new Cisco APIC to the cluster.
- Beginning with Cisco APIC release 6.0(2), an optional argument (*force*) is added to the **decommission** command to allow forcing the decommission operation. The revised command is **decommission controller** *controller-id* [**force**], with the following behaviors:
- When *force* is not declared, the decommission proceeds only if the cluster is not in an unhealthy or upgrade state, where a decommission may not be proper.
 - When *force* is declared, the decommission proceeds regardless of the cluster state.
- For example, `decommission controller 3 force` decommissions APIC3 regardless of the cluster state.
- Step 4** To commission the new Cisco APIC, follow these steps:
- a) Disconnect the old Cisco APIC from the fabric.
 - b) Connect the replacement Cisco APIC to the fabric.
- The new Cisco APIC appears in the Cisco APIC GUI menu **System > Controllers > apic_controller_name > Cluster as Seen by Node** in the **Unauthorized Controllers** list.
- c) Commission the new Cisco APIC using the **controller** *controller-id* **commission** command.
 - d) Boot the new Cisco APIC.
 - e) Allow several minutes for the new Cisco APIC information to propagate to the rest of the cluster.
- The new Cisco APIC appears in the Cisco APIC GUI menu **System > Controllers > apic_controller_name > Cluster as Seen by Node** in the **Active Controllers** list.
-

What to do next

For each decommissioned controller, verify that the operational state of the controller is unregistered and that the controller is no longer in service in the cluster.



Note If a decommissioned Cisco APIC is not promptly removed from the fabric, it might be rediscovered, which could cause problems. In that case, follow the instructions in [Reducing the APIC Cluster Size, on page 244](#) to remove the controller.

Reducing the APIC Cluster Size

Follow these guidelines to reduce the APIC cluster size and decommission the APIC controllers that are removed from the cluster:



Note Failure to follow an orderly process to decommission and power down APIC controllers from a reduced cluster can lead to unpredictable outcomes. Do not allow unrecognized APIC controllers to remain connected to the fabric.

- Reducing the cluster size increases the load on the remaining APIC controllers. Schedule the APIC controller size reduction at a time when the demands of the fabric workload will not be impacted by the cluster synchronization.
- If one or more of the APIC controllers' health status in the cluster is not "fully fit", remedy that situation before proceeding.
- Reduce the cluster target size to the new lower value. For example if the existing cluster size is 6 and you will remove 3 controllers, reduce the cluster target size to 3.
- Starting with the highest numbered controller ID in the existing cluster, decommission, power down, and disconnect the APIC controller one by one until the cluster reaches the new lower target size.

Upon the decommissioning and removal of each controller, the APIC synchronizes the cluster.



Note After decommissioning an APIC controller from the cluster, power it down and disconnect it from fabric. Before returning it to service, do a wiped clean back to factory reset.

- Cluster synchronization stops if an existing APIC controller becomes unavailable. Resolve this issue before attempting to proceed with the cluster synchronization.
- Depending on the amount of data the APIC must synchronize upon the removal of a controller, the time required to decommission and complete cluster synchronization for each controller could be more than 10 minutes per controller.



Note Complete the entire necessary decommissioning steps, allowing the APIC to complete the cluster synchronization accordingly before making additional changes to the cluster.

Contracting the Cisco APIC Cluster

Contracting the Cisco APIC cluster is the operation to decrease any size mismatches, from a cluster size of N to size N - 1, within legal boundaries. As the contraction results in increased computational and memory

load for the remaining APICs in the cluster, the decommissioned APIC cluster slot becomes unavailable by operator input only.

During cluster contraction, you must begin decommissioning the last APIC in the cluster first and work your way sequentially in reverse order. For example, APIC4 must be decommissioned before APIC3, and APIC3 must be decommissioned before APIC2.

Switching Over Active APIC with Standby APIC Using CLI

Use this procedure to switch over an active APIC with a standby APIC.

Procedure

Step 1 `replace-controller replace ID number Backup serial number`

Replaces an active APIC with an standby APIC.

Example:

```
apic1#replace-controller replace 2 FCH1804V27L
Do you want to replace APIC 2 with a backup? (Y/n): Y
```

Step 2 `replace-controller reset ID number`

Resets fail over status of the active controller.

Example:

```
apic1# replace-controller reset 2
Do you want to reset failover status of APIC 2? (Y/n): Y
```

Verifying Cold Standby Status Using the CLI

Procedure

To verify the Cold Standby status of APIC, log in to the APIC as admin and enter the command **show controller**.

```
apic1# show controller
Fabric Name       : vegas
Operational Size  : 3
Cluster Size     : 3
Time Difference   : 496
Fabric Security Mode : strict
```

ID	Pod	Address	In-Band IPv4	In-Band IPv6	OOB IPv4
OOB IPv6			Version	Flags Serial Number	Health
----	----	-----	-----	-----	-----
----	----	-----	-----	-----	-----

```

1* 1 10.0.0.1 0.0.0.0 fc00::1 172.23.142.4
fe80::26e9:b3ff:fe91:c4e0 2.2(0.172) crva- FCH1748V0DF fully-fit
2 1 10.0.0.2 0.0.0.0 fc00::1 172.23.142.6
fe80::26e9:bf8f:fe91:f37c 2.2(0.172) crva- FCH1747V0YF fully-fit
3 1 10.0.0.3 0.0.0.0 fc00::1 172.23.142.8
fe80::4e00:82ff:fead:bc66 2.2(0.172) crva- FCH1725V2DK fully-fit
21~ 10.0.0.21 ----- FCH1734V2DG

```

Flags - c:Commissioned | r:Registered | v:Valid Certificate | a:Approved | f/s:Failover fail/success
 (*)Current (~)Standby

Registering an Unregistered Switch Using the CLI

Use this procedure to register a switch from the **Nodes Pending Registration** tab on the **Fabric Membership** work pane using the CLI.



Note This procedure is identical to "Adding a Switch Before Discovery Using the CLI". When you execute the command, the system determines if the node exists and, if not, adds it. If the node exists, the system registers it.

Procedure

	Command or Action	Purpose
Step 1	[no] system switch-id serial-number switch-id name pod id role leaf node-type tier-2-leaf	Adds the switch to the pending registration list.

Adding a Switch Before Discovery Using the CLI

Use this procedure to add a switch to the **Nodes Pending Registration** tab on the **Fabric Membership** work pane using the CLI.



Note This procedure is identical to "Registering an Unregistered Switch Using the CLI". When you execute the command, the system determines if the node exists and, if not, adds it. If the node does exist, the system registers it.

Procedure

```
[no] system switch-id serial-number switch-id name pod id role leaf node-type tier-2-leaf
```

Adds the switch to the pending registration list.

Removing a Switch to Maintenance Mode Using the CLI

Use this procedure to remove a switch to maintenance mode using the CLI.



Note While the switch is in maintenance mode, CLI 'show' commands on the switch show the front panel ports as being in the up state and the BGP protocol as up and running. The interfaces are actually shut and all other adjacencies for BGP are brought down, but the displayed active states allow for debugging.

Procedure

```
[no]debug-switch node_id or node_name
```

Removes the switch to maintenance mode.

Inserting a Switch to Operation Mode Using the CLI

Use this procedure to insert a switch to operational mode using the CLI.

Procedure

```
[no]no debug-switch node_id or node_name
```

Inserts the switch to operational mode.

Configuring a Remote Location Using the NX-OS Style CLI

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

Procedure

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

Examples

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

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

Finding Your Switch Inventory Using the NX-OS CLI

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

Procedure

Find your switch inventory as follows:

Example:

```

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

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

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

  Device name: ifav41-leaf204
  bootflash:   62522368 kB

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

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

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

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

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

Module1 ok
  Module type is : 48x10/25G

```

```

1 submodules are present
Model number is N9K-C93180YC-EX
H/W version is 0.2110
Part Number is 73-17776-02
Part Revision is 11
Manufacture Date is Year 20 Week 10
Serial number is FDO20101H1W
CLEI code is 73-17776-02

```

GEM ok

```

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

```

```

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

```

PS1 shut

```

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

```

PS2 ok

```

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

```

```

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

```

FT1 ok

```

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

```

FT2 ok

```

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

```

FT3 ok

```

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

```

FT4 ok

```
Fan4(sys_fan4) (fan_model:NXA-FAN-30CFM-F) is inserted  
but info is not available
```

=====

Verifying the Cisco APIC Cluster Using the CLI

Cisco Application Policy Infrastructure Controller (APIC) release 4.2.(1) introduces the **cluster_health** command, which enables you to verify the Cisco APIC cluster status step-by-step. The following output example demonstrates a scenario where everything is fine except for one node (ID 1002), which is inactive.



Note To use the **cluster_health** command, you must be logged in as admin.

Procedure

To verify the cluster status:

```
F1-APIC1# cluster_health  
Password:  
  
Running...  
  
Checking Wiring and UUID: OK  
Checking AD Processes: Running  
Checking All Apics in Commission State: OK  
Checking All Apics in Active State: OK  
Checking Fabric Nodes: Inactive switches: ID=1002 (IP=10.1.176.66/32)  
Checking Apic Fully-Fit: OK  
Checking Shard Convergence: OK  
Checking Leadership Degration: Optimal leader for all shards  
Ping OOB IPs:  
APIC-1: 172.31.184.12 - OK  
APIC-2: 172.31.184.13 - OK  
APIC-3: 172.31.184.14 - OK  
Ping Infra IPs:  
APIC-1: 10.1.0.1 - OK  
APIC-2: 10.1.0.2 - OK  
APIC-3: 10.1.0.3 - OK  
Checking APIC Versions: Same (4.2(0.261a))  
Checking SSL: OK  
  
Done!
```

Table 8: Cluster_Health Verification Steps

Step	Description
Checking Wiring and UUID	<p>Leaf switches provide infra connectivity between each Cisco APIC by detecting the Cisco APICs using LLDP. This step checks wiring issues between a leaf and a Cisco APIC that is detected during LLDP discovery.</p> <p>Any issues in here implies a leaf switch cannot provide infra connectivity for a Cisco APIC as it doesn't have a valid information. For example, a Cisco APIC UUID mismatch means the new APIC2 has a different UUID than the previously known APIC2.</p> <p>UUID – Universally Unique ID, or chassis ID in some outputs</p>
Checking AD Processes	Cisco APIC clustering is handled by the Appliance Director process on each Cisco APIC. This step checks if the process is running correctly.
Checking All APICs in Commission State	To complete the Cisco APIC clustering, all Cisco APICs need to be commissioned.
Checking All APICs in Active State	To complete the Cisco APIC clustering, all commissioned Cisco APICs need to be active. If it is not active, the Cisco APIC may not be up yet.
Checking Fabric Nodes: Inactive switches	The Cisco APIC's communication are through infra connectivity provided by leaf and spine switches. This step checks inactive switches to ensure switches are providing infra connectivity.
Checking APIC Fully-Fit	When Cisco APICs have established IP reachability to each other through infra network, it will synchronize its database to each other. When the synchronization completes, the status of all Cisco APICs become "Fully-Fit." Otherwise, the status will be "Data Layer Partially Diverged," and so on.
Checking Shard Convergence	When Cisco APICs are not fully-fit, database shards need to be checked to see which service is not fully synchronized. If there is any service that has problems in synchronization, you may reach out to Cisco TAC for further troubleshooting.
Checking Leadership Degradation	In ACI, each database shard has one leader shard distributed to each Cisco APIC in the cluster. This step shows if all shards have an optimal leader. If there is an issue in here when all Cisco APICs are up, you may reach out to Cisco TAC for further troubleshooting.

Step	Description
Ping OOB IPs	This step is to check if all Cisco APICs are up and operational by pinging the OOB IP which is configured separately from clustering.
Ping Infra IPs	This step is to check if there is infra connectivity between each Cisco APIC. Cisco APIC clustering is performed through infra connectivity instead of OOB.
Checking APIC Versions	All Cisco APICs should be on a same version to complete clustering.
Checking SSL	All Cisco APICs need to have a valid SSL that should be built-in when a Cisco APIC is shipped as an appliance. Without a valid SSL, the server cannot operate the Cisco APIC OS correctly.



APPENDIX **B**

Configuring the Cisco APIC Using the REST API

- [Expanding the APIC Cluster Using the REST API, on page 255](#)
- [Contracting the APIC Cluster Using the REST API, on page 256](#)
- [Reducing the APIC Cluster Size, on page 258](#)
- [Switching Over Active APIC with Standby APIC Using REST API, on page 258](#)
- [Registering an Unregistered Switch Using the REST API, on page 259](#)
- [Adding a Switch Before Discovery Using the REST API, on page 260](#)
- [Removing a Switch to Maintenance Mode Using the REST API, on page 260](#)
- [Inserting a Switch to Operational Mode Using the REST API, on page 261](#)
- [Configuring a Remote Location Using the REST API, on page 261](#)
- [Sending an On-Demand Tech Support File Using the REST API, on page 262](#)
- [Finding Your Switch Inventory Using the REST API, on page 262](#)

Expanding the APIC Cluster Using the REST API

The cluster drives its actual size to the target size. If the target size is higher than the actual size, the cluster size expands.

Procedure

Step 1 Set the target cluster size to expand the APIC cluster size.

Example:

```
POST
https://<IP address>/api/node/mo/uni/controller.xml
<infraClusterPol name='default' size=3/>
```

Step 2 Physically connect the APIC controllers that you want to add to the cluster.

Contracting the APIC Cluster Using the REST API

Use this procedure to shrink the cluster size by removing controllers. For more information about contracting the cluster size, see [Contracting the Cisco APIC Cluster](#), on page 244.



Note Beginning with Cisco APIC Release 6.0(2), two additional properties are added to the API command to allow forcing the decommission operation. The new object properties are:

- `infraClusterPol:shrink`
 - `false`: (default) If the target cluster size (`infraClusterPol:size`) is less than the current operational cluster size, the APICs to be removed must be decommissioned manually, as in earlier releases.
 - `true`: If the target cluster size is less than the current operational cluster size, a cluster shrink decommission is triggered. The APICs to be removed are decommissioned automatically, beginning with the APIC with the highest controller ID number.
- `infraWiNode:force`
 - `false`: (default) The decommission proceeds only if the cluster is not in an unhealthy or upgrade state, where a decommission may not be proper.
 - `true`: The decommission proceeds regardless of the cluster state.

This example shows how to contract the cluster from three APIC controllers to a single controller. To achieve a target size of one, APIC3 and APIC2 must be decommissioned, in that order.

Procedure

Step 1 Set the target cluster size so as to contract the APIC cluster size.

When the cluster size is reduced with `shrink='true'`, the APICs to be removed are decommissioned automatically. Otherwise, they must be decommissioned manually.

Example:

Cisco APIC Release 6.0(1) and earlier:

```
POST
https://<IP address>/api/node/mo/uni/controller.xml
<infraClusterPol name='default' size=1 />
```

You must now manually decommission the APICs to be removed, as shown in the next steps.

Example:

Cisco APIC Release 6.0(2) and later, using 'shrink' property:

```
POST
https://<IP address>/api/node/mo/uni/controller.xml
<infraClusterPol name='default' size=1 shrink='true' />
```


With `shrink='true'`, the following steps can be skipped. The APICs to be removed are decommissioned automatically.

```
POST
https://<IP address>/api/node/mo/uni/controller.xml
<infraClusterPol name='default' size=1 shrink='false' />
```

With `shrink='false'`, you must now manually decommission the APICs to be removed, as shown in the next steps.

Step 2 Decommission APIC3 on APIC1 for cluster contraction.

Example:

Cisco APIC Release 6.0(1) and earlier:

```
POST
https://<IP address>/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=3 adminSt='out-of-service' />
```

The decommission proceeds only if the cluster is in a healthy state.

Example:

Cisco APIC Release 6.0(2) and later, using 'force' property:

```
POST
https://<IP address>/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=3 adminSt='out-of-service' force='true' />
```

With `force='true'`, the decommission proceeds regardless of the cluster state.

Step 3 Decommission APIC2 on APIC1 for cluster contraction.

Example:

Cisco APIC Release 6.0(1) and earlier:

```
POST
https://<IP address>/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=2 adminSt='out-of-service' />
```

The decommission proceeds only if the cluster is in a healthy state.

Example:

Cisco APIC Release 6.0(2) and later, using 'force' property:

```
POST
https://<IP address>/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=2 adminSt='out-of-service' force='false' />
```

With `force='false'`, the decommission proceeds only if the cluster is in a healthy state.

The operation cluster size shrinks only after the last appliance is decommissioned, and not after the administrative size is changed. Verify after each controller is decommissioned that the operational state of the controller is unregistered and that the controller is no longer in service in the cluster.



Note If a decommissioned APIC controller is not promptly removed from the fabric, it might be rediscovered, which could cause problems. In that case, follow the instructions in [Reducing the APIC Cluster Size, on page 244](#) to remove the controller.

Reducing the APIC Cluster Size

Follow these guidelines to reduce the APIC cluster size and decommission the APIC controllers that are removed from the cluster:



Note Failure to follow an orderly process to decommission and power down APIC controllers from a reduced cluster can lead to unpredictable outcomes. Do not allow unrecognized APIC controllers to remain connected to the fabric.

- Reducing the cluster size increases the load on the remaining APIC controllers. Schedule the APIC controller size reduction at a time when the demands of the fabric workload will not be impacted by the cluster synchronization.
- If one or more of the APIC controllers' health status in the cluster is not "fully fit", remedy that situation before proceeding.
- Reduce the cluster target size to the new lower value. For example if the existing cluster size is 6 and you will remove 3 controllers, reduce the cluster target size to 3.
- Starting with the highest numbered controller ID in the existing cluster, decommission, power down, and disconnect the APIC controller one by one until the cluster reaches the new lower target size.

Upon the decommissioning and removal of each controller, the APIC synchronizes the cluster.



Note After decommissioning an APIC controller from the cluster, power it down and disconnect it from fabric. Before returning it to service, do a wiped clean back to factory reset.

- Cluster synchronization stops if an existing APIC controller becomes unavailable. Resolve this issue before attempting to proceed with the cluster synchronization.
- Depending on the amount of data the APIC must synchronize upon the removal of a controller, the time required to decommission and complete cluster synchronization for each controller could be more than 10 minutes per controller.



Note Complete the entire necessary decommissioning steps, allowing the APIC to complete the cluster synchronization accordingly before making additional changes to the cluster.

Switching Over Active APIC with Standby APIC Using REST API

Use this procedure to switch over an active APIC with standby APIC using REST API.

Procedure

Switch over active APIC with standby APIC.

```
URL for POST: https://ip
address/api/node/mo/topology/pod-initiator_pod_id/node-initiator_id/av.xml
Body: <infraWiNode id=outgoing_apic_id targetMbSn=backup-serial-number/>
where initiator_id = id of an active APIC other than the APIC being replaced.
pod-initiator_pod_id = pod ID of the active APIC
backup-serial-number = serial number of standby APIC
```

Example:

```
https://ip address/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=2 targetMbSn=FCH1750V00Q/>
```

Registering an Unregistered Switch Using the REST API

Use this procedure to register a switch from the **Nodes Pending Registration** tab on the **Fabric Membership** work pane using the REST API.



Note This procedure is identical to "Adding a Switch Before Discovery Using the REST API". When you apply the code, the system determines if the node exists and, if not, adds it. If the node does exist, the system registers it.

Procedure

Add a switch description.

Example:

```
POST
https://<IP address>/api/policymgr/mo/uni.xml

<!-- /api/policymgr/mo/uni.xml -->
<polUni>
<ctrlrInst>
  <fabricNodeIdentPol>
    <fabricNodeIdentP nodeType="tier-2-leaf" podId="1" serial="XXXXXXXX"
      name="tier-2-leaf-leaf1" nodeId="101"/>
  </fabricNodeIdentPol>
</ctrlrInst>
</polUni>
```

Adding a Switch Before Discovery Using the REST API

Use this procedure to add a switch to the **Nodes Pending Registration** tab on the **Fabric Membership** work pane using the REST API.



Note This procedure is identical to "Registering an Unregistered Switch Using the REST API". When you apply the code, the system determines if the node exists and, if not, adds it. If the node does exist, the system registers it.

Procedure

Add a switch description.

Example:

```
POST
https://<IP address>/api/policymgr/mo/uni.xml

<!-- /api/policymgr/mo/uni.xml -->
<polUni>
<ctrlrInst>
  <fabricNodeIdentPol>
    <fabricNodeIdentP  nodeType="tier-2-leaf" podId="1" serial="XXXXXXXXX"
      name="tier-2-leaf1" nodeId="101"/>
  </fabricNodeIdentPol>
</ctrlrInst>
</polUni>
```

Removing a Switch to Maintenance Mode Using the REST API

Use this procedure to remove a switch to maintenance mode using the REST API.

Procedure

Remove a switch to maintenance mode.

Example:

```
POST
https://<IP address>/api/node/mo/uni/fabric/outofsvc.xml

<fabricOOServicePol
  descr=""
  dn=""
  name="default"
```

```
    nameAlias=""
    ownerKey=""
    ownerTag="">
<fabricRsDecommissionNode
  debug="yes"
  dn=""
  removeFromController="no"
  tDn="topology/pod-1/node-102"/>
</fabricOOServicePol>
```

Inserting a Switch to Operational Mode Using the REST API

Use this procedure to insert a switch to operational mode using the REST API.

Procedure

Insert a switch to operational mode.

Example:

```
POST
https://<IP address>/api/node/mo/uni/fabric/outofsvc.xml

<fabricOOServicePol
  descr=""
  dn=""
  name="default"
  nameAlias=""
  ownerKey=""
  ownerTag="">
<fabricRsDecommissionNode
  debug="yes"
  dn=""
  removeFromController="no"
  tDn="topology/pod-1/node-102"
  status="deleted"/>
</fabricOOServicePol>
```

Configuring a Remote Location Using the REST API

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

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

Sending an On-Demand Tech Support File Using the REST API

Procedure

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

Example:

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

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

</fileRemotePath>
```

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

Example:

```
<dbgexpTechSupOnD upgradeLogs="no" startTime="unspecified" name="Tech_Support_9-20-16"
  exportToController="no" endTime="unspecified" dn="uni/fabric/tsod-Tech_Support_9-20-16"
  descr=""
  compression="gzip" category="forwarding" adminSt="untriggered">
  <dbgexpRsExportDest tDn="uni/fabric/path-ToSupport"/>
  <dbgexpRsTsSrc tDn="topology/pod-1/node-102/sys"/>
  <dbgexpRsTsSrc tDn="topology/pod-1/node-103/sys"/>
  <dbgexpRsTsSrc tDn="topology/pod-1/node-101/sys"/>
  <dbgexpRsData tDn="uni/fabric/tscont"/>
</dbgexpTechSupOnD>
<fabricFuncP>
  <fabricCtrlrPGrp name="default">
    <fabricRsApplTechSupOnDemand tnDbgexpTechSupOnDName=" Tech_Support_9-20-16"/>
  </fabricCtrlrPGrp>
</fabricFuncP>
```

Finding Your Switch Inventory Using the REST API

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

Procedure

Find your switch inventory as follows:

Example:

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

The following response is returned:

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

