



Cisco Cloud APIC for AWS User Guide, Release 25.0(1)-25.0(4)

First Published: 2021-09-20

Last Modified: 2022-12-30

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 Cloud APIC 5
	Overview 5
	External Network Connectivity 6
	Understanding Supported Routing and Security Policies 7
	Routing and Security Policies: Releases Prior to 25.0(1) 7
	Routing and Security Policies: Release 25.0(1) 7
	Routing Policies: Release 25.0(2) 9
	Source Interface Selection for Tunnels 11
	General Guidelines and Limitations for Cisco Cloud APIC 11
	About the Cisco Cloud APIC GUI 14
	Understanding the Cisco Cloud APIC GUI Icons 15

CHAPTER 3	Cisco Cloud APIC Policy Model 19
	About the ACI Policy Model 19
	Policy Model Key Characteristics 19
	Logical Constructs 20
	The Cisco ACI Policy Management Information Model 21
	Tenants 23
	Cloud Context Profile 24
	CCR 24
	About the Cisco Catalyst 8000V 24

- Private IP Address Support for Cisco Cloud APIC and CCR in AWS 27
- Communicating to External Sites From Regions Without a CCR 28
- Support for ECMP Forwarding from Remote Sites for CCRs 32
- Preference For Routes to CCRs in Regions with Local CIDRs 32
- Availability Zones 32
 - Migrating from Virtual Availability Zones to Cloud Availability Zones 33
 - Guidelines and Limitations 34
- VRFs 34
- Cloud Application Profiles 35
- Cloud Endpoint Groups 35
- Contracts 37
 - Filters and Subjects Govern Cloud EPG Communications 38
- About the Cloud Template 39
- Managed Object Relations and Policy Resolution 42
- Default Policies 43
- Shared Services 44

CHAPTER 4

- Configuring Cisco Cloud APIC Components 45**
 - About Configuring the Cisco Cloud APIC 45
 - Configuring the Cisco Cloud APIC Using the GUI 45
 - Creating a Tenant Using the Cisco Cloud APIC GUI For Release 4.2(2) and Earlier 45
 - Creating a Tenant Using the Cisco Cloud APIC GUI For Release 4.2(3) and Later 46
 - Configure a Tenant AWS Provider For Release 4.2(2) and Earlier 47
 - Configuring a Tenant AWS Provider For Release 4.2(3) and Later 50
 - Creating an Application Profile Using the Cisco Cloud APIC GUI 55
 - Creating a VRF Using the Cisco Cloud APIC GUI 55
 - Creating an External Network Using the Cisco Cloud APIC GUI 57
 - Configuring the Global Inter-VRF Route Leak Policy 61
 - Configuring Leak Routes Using the Cisco Cloud APIC GUI 62
 - Configuring Inter-VRF Route Leaking Using the Cisco Cloud APIC GUI 62
 - Configuring Leak Routes for Internal VRFs Using the Cisco Cloud APIC GUI 65
 - Enabling Connectivity From the AWS Site to External Devices 67
 - Downloading the External Device Configuration Files 67
 - Enabling Connectivity From the AWS Site to External Devices 67

Creating an EPG Using the Cisco Cloud APIC GUI	71
Creating a Contract Using the Cisco Cloud APIC GUI	76
Specifying Consumer and Provider EPGs Using the Cisco Cloud APIC	78
Creating a Filter Using the Cisco Cloud APIC GUI	79
Creating a Cloud Context Profile Using the Cisco Cloud APIC GUI	80
Configuring Instances in AWS	82
Creating a Backup Configuration Using the Cisco Cloud APIC GUI	84
Creating a Tech Support Policy Using the Cisco Cloud APIC GUI	87
Creating a Trigger Scheduler Using the Cisco Cloud APIC GUI	88
Creating a Remote Location Using the Cisco Cloud APIC GUI	90
Creating a Login Domain Using the Cisco Cloud APIC GUI	91
Creating a Provider Using the Cisco Cloud APIC GUI	94
Creating a Security Domain Using the Cisco Cloud APIC GUI	98
Creating a Role Using the Cisco Cloud APIC GUI	99
Creating an RBAC Rule Using the Cisco Cloud APIC GUI	104
Creating a Certificate Authority Using the Cisco Cloud APIC GUI	105
Creating a Key Ring Using the Cisco Cloud APIC GUI	106
Creating a Local User Using the Cisco Cloud APIC GUI	108
Managing Regions (Configuring a Cloud Template) Using the Cisco Cloud APIC GUI	110
Configuring Cisco Cloud APIC Using the REST API	112
Creating a Tenant Using the REST API	112
Creating a Contract Using the REST API	113
Creating a Cloud Context Profile Using the REST API	113
Managing a Cloud Region Using the REST API	115
Creating a Filter Using the REST API	116
Creating an Application Profile Using the REST API	116
Creating a Cloud EPG Using the REST API	117
Creating an External Cloud EPG Using the REST API	118
Creating a Cloud Template Using the REST API	118
Configuring VRF Leak Routes Using the REST API	120
Configuring the Source Interface Selection for Tunnels Using the REST API	122

CHAPTER 5
Viewing System Details 123

Monitoring VM Host Metrics	123
----------------------------	-----

- Monitoring VM Host Metrics Using the GUI 123
- Monitoring VM Host Metrics Using the REST API 125
- Viewing Application Management Details 126
- Viewing Cloud Resource Details 127
- Viewing Operations Details 128
- Viewing Infrastructure Details 130
- Viewing Administrative Details 130
- Viewing Health Details Using the Cisco Cloud APIC GUI 132

CHAPTER 6

Deploying Layer 4 to Layer 7 Services 135

- Overview 135
 - About Application Load Balancers 135
 - Dynamic Server Attachment to Server Pool 137
 - About Service Graphs 138
 - About Function Nodes 139
 - About Terminal Nodes 139
- Deploying a Service Graph 139
 - Deploying the Service Graph Using the Cloud APIC GUI 139
 - Creating a Load Balancer Using the Cisco Cloud APIC GUI 139
 - Creating a Service Graph Template Using the Cisco Cloud APIC GUI 141
 - Deploying Layer 4 to Layer 7 Services Using the Cisco Cloud APIC GUI 142
 - Deploying a Service Graph Using the REST API 147
 - Creating an Internal-Facing Load Balancer Using the REST API 147
 - Configuring an Internet-Facing Load Balancer Using the REST API 147
 - Creating a Service Graph Using the REST API 148
 - Attaching a Service Graph Using the REST API 148
 - Configuring an HTTP Service Policy Using the REST API 149
 - Configuring a Key Ring Using the REST API 149
 - Creating an HTTPS Service Policy Using the REST API 151

CHAPTER 7

Cisco Cloud APIC Statistics 153

- About Cisco Cloud APIC Statistics 153
- AWS Networking Interface Statistics Collection 153
- Cisco Cloud APIC Endpoints and cloudEPg Statistics Processing 154

Cisco Cloud APIC Statistics Filters	154
AWS Transit Gateway Statistics	155
Enabling VPC Flow Logs	156
Enabling VPC Flow Logs Using the Cisco Cloud APIC GUI	156
Enabling VPC Flow Logs Using the REST API	158
Cloud Router Statistics	159

CHAPTER 8**Cisco Cloud APIC Security 163**

Access, Authentication, and Accounting	163
Configuration	163
Configuring TACACS+, RADIUS, LDAP and SAML Access	164
Overview	164
Configuring Cloud APIC for TACACS+ Access	164
Configuring Cloud APIC for RADIUS Access	165
Configuring a Cisco Secure Access Control Server for RADIUS and TACACS+ Access to the Cloud APIC	167
Configuring LDAP Access	167
Configuring Windows Server 2008 LDAP for APIC Access with Cisco AVPair	167
Configuring Cloud APIC for LDAP Access	167
Configuring Cloud APIC for SAML Access	169
About SAML	169
Configuring Cloud APIC for SAML Access	170
Setting Up a SAML Application in Okta	171
Setting Up a Relying Party Trust in AD FS	171
Configuring HTTPS Access	171
About HTTPS Access	172
Guidelines for Configuring Custom Certificates	172
Configuring a Custom Certificate for Cisco ACI HTTPS Access Using the GUI	172

CHAPTER 9**Configuration Drifts 175**

Configuration Drift Notifications and Faults	175
Accessing the Main Configuration Drift Page	176
Checking for Missing Contracts Configuration	179
Checking for Missing EPGs Configuration	180

Checking for Missing VRFs Configuration 182

Configuration Drift Troubleshooting 183

CHAPTER 10

AWS Transit Gateway on Cisco Cloud APIC 185

AWS Transit Gateway on Cisco Cloud APIC 185

APPENDIX A

Cisco Cloud APIC Error Codes 187

Cisco Cloud APIC Error Codes 187



CHAPTER 1

New and Changed Information

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

New and Changed Information

The following table provides an overview of the significant changes to the organization and features in 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 Behavior in Cisco Cloud APIC for Release 25.0(4)

Feature or Change	Description	Where Documented
Support for PAYG Licensing Model on Cisco Catalyst 8000V in Cisco Cloud APIC	Cisco Cloud APIC supports Pay-As-You-Go (PAYG) Licensing Model on Cisco Catalyst 8000V which allows users to deploy a Catalyst 8000V instance in the cloud based on the VM size and purchase the usage on an hourly basis.	

Table 2: New Features and Changed Behavior in Cisco Cloud APIC for Release 25.0(3)

Feature or Change	Description	Where Documented
Move from the Cisco Cloud Services Router 1000v to the Cisco Catalyst 8000V	Cisco Cloud APIC moves from the Cisco Cloud Services Router 1000v to the Cisco Catalyst 8000V beginning with release 25.0(3).	

Feature or Change	Description	Where Documented
Terms used for Cisco Cloud Services Router 1000v and Cisco Catalyst 8000V	<p>The following terms are used for the two types of routers described above:</p> <ul style="list-style-type: none"> • CSR: Short for Cloud Services Router. Refers to the Cisco Cloud Services Router 1000v, used in releases prior to release 25.0(3). • CCR: Short for Cisco Cloud Router. Refers to the Cisco Catalyst 8000V, used in release 25.0(3) and later. <p>In addition, throughout this document, CCR is used as a generic term for either of the routers described above, depending on your release.</p>	
Change in name of Multi-Site Orchestrator	Cisco ACI Multi-SiteOrchestrator (MSO) has changed to Cisco Nexus Dashboard Orchestrator (NDO) beginning with the MSO release 3.4.1 on August 15, 2021. Every instance of MSO is now NDO in this Cisco Cloud APIC documentation.	

Table 3: New Features and Changed Behavior in Cisco Cloud APIC for Release 25.0(2)

Feature or Change	Description	Where Documented
Support for multiple (greater than two) availability zones in AWS for Cisco Cloud APIC	Support is now provided for multiple (greater than two) availability zones in AWS for Cisco Cloud APIC	Cisco Cloud APIC Policy Model, on page 19
Support for configuring routing and security policies independently in AWS	<p>Beginning with release 25.0(2), the following updates are available for the routing policies:</p> <ul style="list-style-type: none"> • Support for route maps-based route leaking between a pair of internal VRFs • Support for the internal VRF route leak policy, which allows you to choose whether you want to use contract-based routing or maps-based routing between a pair of internal VRFs 	<ul style="list-style-type: none"> • About Cisco Cloud APIC, on page 5 • Configuring Cisco Cloud APIC Components, on page 45

Feature or Change	Description	Where Documented
CCR IPsec tunnels can now use any of the three available data interfaces for external branch connectivity	<p>Prior to release 25.0(2), all the tunnels to external networks are originated from one specific interface on the CCR router (the GigabitEthernet3 interface, or cloudHostIfp-2).</p> <p>Beginning with release 25.0(2), support is now extended where tunnels to the same destination can be formed from the GigabitEthernet2, GigabitEthernet3, and GigabitEthernet4 interfaces. This is supported for tunnels with IKEv2 configurations only.</p>	<ul style="list-style-type: none"> • About Cisco Cloud APIC, on page 5 • Configuring Cisco Cloud APIC Components, on page 45
Support for increased number of cloud regions for workload deployment	Prior to release 25.0(2), you can have a maximum of four regions per site. Beginning with release 25.0(2), you can have a maximum of sixteen regions per site.	

Table 4: New Features and Changed Behavior in Cisco Cloud APIC for Release 25.0(1)

Feature or Change	Description	Where Documented
Change in release numbering for Cisco Cloud APIC	<p>Beginning with release 25.0(1), the release numbering has changed for Cisco Cloud APIC. The sequential order of releases for Cisco Cloud APIC is as follows:</p> <ul style="list-style-type: none"> • 4.1(x) (support for AWS only) • 4.2(x) • 5.0(x) • 5.1(x) • 5.2(x) • 25.0(x) 	
Support for Prometheus Node Exporter on Cisco Cloud APIC	The Prometheus Node Exporter is supported on Cisco Cloud APIC beginning with release 25.0(1).	Viewing System Details, on page 123

Feature or Change	Description	Where Documented
Support for IPv4 connectivity from the infra VPC CCRs to any external device with IPSec/BGP.	Support is now available for IPv4 connectivity from the infra VPC CCRs to any external device with IPSec/BGP.	External Network Connectivity, on page 6
Support for configuring routing policies separately, independent of security policies, between internal and external VRFs when configuring for external connectivity.	Support is now available for configuring routing policies separately, independent of security policies, between internal and external VRFs when configuring for external connectivity.	Understanding Supported Routing and Security Policies, on page 7



CHAPTER 2

About Cisco Cloud APIC

- [Overview, on page 5](#)
- [External Network Connectivity, on page 6](#)
- [Understanding Supported Routing and Security Policies, on page 7](#)
- [Source Interface Selection for Tunnels, on page 11](#)
- [General Guidelines and Limitations for Cisco Cloud APIC, on page 11](#)
- [About the Cisco Cloud APIC GUI, on page 14](#)

Overview

Cisco Application Policy Infrastructure Controller (APIC) Release 4.1(1) introduces Cisco Cloud APIC, which is a software deployment of Cisco APIC that you deploy on a cloud-based virtual machine (VM). When deployed, Cisco Cloud APIC:

- Provides an interface that is similar to the existing Cisco APIC to interact with the AWS public cloud
- Automates the deployment and configuration of cloud constructs
- Configures the cloud router control plane
- Configures the data path between the on-premises Cisco ACI fabric and the cloud site
- Translates Cisco ACI policies to cloud native construct
- Discovers endpoints
- Provides a consistent policy, security, and analytics for workloads deployed either on or across on-premises data centers and the public cloud



Note

- Cisco Multi-Site pushes the MP-BGP EVPN configuration to the on-premises spine switches
 - On-premises VPN routers require a manual configuration for IPsec
-

- Provides an automated connection between on-premises data centers and the public cloud with easy provisioning and monitoring

- Policies are pushed by Cisco Nexus Dashboard Orchestrator to the on-premises and cloud sites, and Cisco Cloud APIC translates the policies to the cloud to keep the policies consistent with the on-premises site

For more information about extending Cisco ACI to the public cloud, see the *Cisco Cloud APIC Installation Guide*.

When the Cisco Cloud APIC is up and running, you can begin adding and configuring Cisco Cloud APIC components. This document describes the Cisco Cloud APIC policy model and explains how to manage (add, configure, view, and delete) the Cisco Cloud APIC components using the GUI and the REST API.

External Network Connectivity

Prior to release 25.0(1), external network connectivity for Cisco Cloud APIC with AWS was available only by using EVPN connectivity from the CCRs in the infra VPC.

Beginning with release 25.0(1), support is also available for IPv4 connectivity from the infra VPC CCRs to any external device with IPSec/BGP. This IPSec/BGP external connectivity allows Cisco Cloud APIC to connect to branch offices.

The following sections provide more information on the components that allow for the new external network connectivity provided in release 25.0(1).

External VRF

An **external VRF** is a unique VRF that does not have any presence in the cloud but is associated with one or more external networks. As opposed to an internal VRF, which is a VRF that is used to host the VPCs and is associated with a cloud context profile, an external VRF is not referred to in any cloud context profile used by Cisco Cloud APIC.

An external VRF represents an external network that is connected to other cloud sites or to on-premises branch offices. Multiple cloud VRFs can leak routes to an external VRF or can get the routes from an external VRF. When an external network is created on an external VRF, inter-VRF routing is set up so that routes received and advertised on the external network are received or advertised on the external VRF.

Connections to Non-ACI External Devices

For release 25.0(1), the existing external connectivity model is extended to provide connectivity from AWS CCRs to any non-ACI external device. IPv4 sessions are created on an external VRF from the infra VPC CCRs to these non-ACI external devices, and inter-VRF routing is set up between the external VRF and the site local VRFs.

Following are the guidelines and limitations for this type of connectivity:

- You cannot use both EVPN and IPv4 IPSec/BGP to connect from the cloud to the same remote site.

Guidelines and Limitations

Instead of manually selecting all the regions, you have to set `allRegion` to true for the external network connectivity starting in release 25.0(2).

Understanding Supported Routing and Security Policies

Routing and security policies are handled differently, depending on the release that is running on your Cisco Cloud APIC.

Routing and Security Policies: Releases Prior to 25.0(1)

Prior to release 25.0(1), routing and security policies are tightly coupled together. To allow communication between two endpoints that are across EPGs, you must configure contracts. These contracts are used for the following:

- **Routing policies:** Policies used to define routes to establish traffic flow.
- **Security policies:** Rules used for security purposes, such as security group rules or network security group rules.

In other words, contracts inherently serve the dual purpose of configuring both security policies and routing policies. This means that tearing down contracts not only tears down the security policies that govern which traffic to allow and which to deny, it also tears down any policies used to route that traffic. Prior to release 25.0(1), there is no way to configure routing policies without also configuring security policies, and vice versa.

Routing and Security Policies: Release 25.0(1)

Beginning with release 25.0(1), support is now available for configuring routing separately, independent of the security policies.



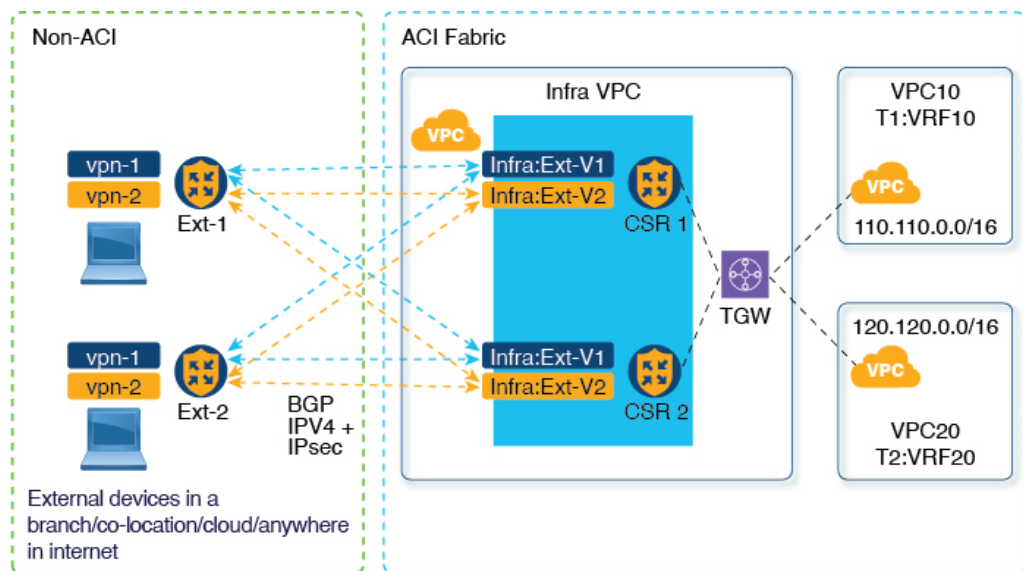
Note The routing and security policies described in this section are specifically for the 25.0(1) release and apply only between internal and external VRFs. For changes in the routing and security policies in the 25.0(2) release, see [Routing Policies: Release 25.0\(2\), on page 9](#).

The procedures for configuring the routing and security policies are here:

- **Routing policy:** You will use the inter-VRF routing feature introduced in release 25.0(1) to configure the routing policy separately. See [Configuring Inter-VRF Route Leaking Using the Cisco Cloud APIC GUI, on page 62](#) for those procedures.
- **Security policy:** After you have configured the routing policy, you will continue to use contracts as you did previously to configure the security policy separately:
 - First create an external EPG. See [Creating an EPG Using the Cisco Cloud APIC GUI, on page 71](#) for those procedures.
 - Then create a contract between the external EPG and the cloud EPG. See [Creating a Contract Using the Cisco Cloud APIC GUI, on page 76](#) for those procedures.

Using inter-VRF routing, you can configure an independent routing policy to specify which routes to leak between a pair of internal and external VRFs when you are setting up routing between a cloud site and a non-ACI site.

The following figure shows an example topology of this sort of configuration. This example topology shows how you can connect to a remote endpoint (vpn-1) behind an external device (Ext-1) which might be located in a non-ACI site. This non-ACI site could be a branch office, co-located or cloud site, or anywhere in the internet that has the capability of BGP IPv4 and IPsec.



In this example, the infra:Ext-V1 is the external VRF on the CCRs in the infra VPC, with BGP IPv4 sessions over IPsec tunnels to the remote devices. The remote endpoint routes are received over these sessions in the infra:Ext-V1 VRF, which are then leaked into the internal VRFs displayed on the right side of the graphic (for example, the T1:VRF10 in VPC10). The reverse leaking routes are also configured.

Route leaking occurs between internal and external VRFs using route maps. Cisco Cloud APIC supports using route maps to configure routing policies independent of security policies only from internal VRFs to external VRFs, and from external VRF to internal VRFs. You will continue to use contracts when configuring routing between a pair of internal VRFs, so routing and security policies are tied together in the configuration process when routing between internal VRFs.

The following list provides more information on situations when you can use **route maps** to configure routing policies independent of security policies, and when you have to use **contracts** where the routing and security policies are tied together.

- Routing situations that use contracts-based routing:
 - Intra-site routing (within and across regions)
 - Inter-site routing (cloud-to-ACI on-premises using EVPN)
 - Cloud-to-cloud routing
 - Route leaking between internal VRFs
- Routing situations that use route map-based routing:
 - Cloud-to-non-ACI on-premises site using L3Out external VRF (no EVPN)
 - Leak specific or all routes from an internal VRF to an external VRF
 - Leak specific or all routes from an external VRF to an internal VRF

Guidelines and Restrictions for Security and Routing Policies in Release 25.0(1)

The following guidelines apply when using inter-VRF routing to leak routes between a pair of VRFs using route maps:

- Routes are always leaked bi-directionally between an internal VRF and the external VRF.
For example, assume there is a user tenant (t1) with an internal VRF (V1) and external VRF (Ext-V1). The route leak must be configured for both of these VRFs bi-directionally.
- You cannot configure "smaller" prefixes to be leaked while a "larger" prefix is already being leaked. For example, configuring the 10.10.10.0/24 prefix will be rejected if you already have the 10.10.0.0/16 prefix configured to be leaked. Similarly, if you configure the 0.0.0.0/0 (leak all) prefix, no other prefix will be allowed to be configured.
- Contracts are not allowed between cloud external EPGs (cloudExtEpgs).
- An external VRF cannot be used for creating cloud EPGs.
- An external VRF always belongs to the infra tenant.
- Leak routing is not supported between external VRFs.

Routing Policies: Release 25.0(2)



Note The routing and security policies described in this section are specifically for the 25.0(2) release. For changes in the routing and security policies in the previous release, see [Routing and Security Policies: Release 25.0\(1\), on page 7](#).

For release 25.0(2), the routing and security policies continue to be split as described in [Routing and Security Policies: Release 25.0\(1\), on page 7](#), but with these additional changes specifically for the routing policies:

- [Route Leaking Between Internal VRFs, on page 9](#)
- [Global Inter-VRF Route Leak Policy, on page 10](#)
- [Guidelines and Limitations, on page 11](#)

Route Leaking Between Internal VRFs

In the previous 25.0(1) release, the inter-VRF route map-based routing feature was introduced, where you can configure an independent routing policy to specify which routes to leak between a pair of internal and external VRFs. This route map-based routing feature applied specifically between internal and external VRFs; when configuring routing between a pair of internal VRFs, you could only use contract-based routing in that situation, as described in [Routing and Security Policies: Release 25.0\(1\), on page 7](#).

Beginning with release 25.0(2), support is now available for route map-based route leaking between a pair of internal VRFs. You will specify how routes are leaked using one of the following options:

- Leak all CIDRS or specific subnet IP addresses associated with the VRF by using:
 - **Leak All** option through the GUI
 - `leakInternalPrefix` field through the REST API

- Leak between a pair of VRFs by using:
 - **Subnet IP** option through the GUI
 - `leakInternalSubnet` field through the REST API

Global Inter-VRF Route Leak Policy

In addition to the support that is now available for route map-based route leaking between a pair of internal VRFs, the internal VRF route leak policy also allows you to choose whether you want to use contract-based routing or route map-based routing between a pair of internal VRFs. This is a global mode configuration available in the First Time Setup to allow a contract-based or route map-based model. Note that when you enable contract-based routing in this global mode, the routes between a pair of internal VRFs can be leaked using contracts only in the absence of route maps.

This policy has the following characteristics:

- This policy is associated with every internal VRF.
- This is a Cisco Cloud APIC-created policy.
- Contract-based routing is disabled by default (turned off) for greenfield cases (when you are configuring a Cisco Cloud APIC for the first time). For upgrades, where you have a Cisco Cloud APIC that was already configured prior to release 25.0(2), contract-based routing is enabled (turned on).

The internal VRF route leak policy is a global policy that is configured in the First Time Setup screen under the infra tenant, where a Boolean flag is used to indicate whether contracts can drive routes in the absence of route maps:

- **Off**: Default setting. Routes are not leaked based on contracts, and are leaked based on route maps instead.
- **On**: Routes are leaked based on contracts in the absence of route maps. When enabled, contracts drive routing when route maps are not configured. When route maps exist, route maps always drives routing.

You can toggle this Boolean flag back and forth. Following are the general recommended steps for toggling this global VRF route leak policy, with more detailed instructions provided in [Configuring Leak Routes for Internal VRFs Using the Cisco Cloud APIC GUI, on page 65](#).

- You should enable contract-based routing in Cisco Cloud APIC for multi-cloud and hybrid-cloud deployments with EVPN.
- For multi-cloud and hybrid-cloud deployments without EVPN, routing is driven through route maps only and not through contracts.
- If you want to disable contract-based routing by toggling from contract-based routing to route map-based routing (toggling to the **Off** setting), this action can be disruptive if route map-based routing is not configured before you've toggled this setting to **Off**.

You should make the following configuration changes before toggling to route map-based routing:

1. Enable route map-based route leaking between all pairs of VRFs that have existing contracts.
2. Disable contract-based routing policy in the global policy.

At that point, you can change the routing policy to route map-based routing, and you can then change the routing to reflect any granularity that is required with the new route map-based routing.

- If you want to enable contract-based routing by toggling from route map-based routing to contract-based routing (toggling to the **On** setting), you do not have to make any configuration changes before toggling to contract-based routing. That's because this setting is an additive operation. In other words, both contract-based and route map-based routing can be enabled between a pair of VRFs. Route maps take precedence over contracts when enabling routing. With route map-based routing enabled, adding contract-based routing should be non-disruptive.

Guidelines and Limitations

The following guidelines and limitations apply for release 25.0(2):

- Routing between external and internal VRFs continues to use route map-based routing only.
- The `leakExternalPrefix` should not overlap with the route to the internet gateway (the external endpoint selector configured for external EPG to perform SSH), otherwise SSH will be broken.

Source Interface Selection for Tunnels

Prior to release 25.0(2), IPsec tunnels to the same destination were not allowed. Beginning with release 25.0(2), you can have more than one tunnel across different external networks to the same destination. This is done in the GUI by using different source interfaces (2,3, or 4) or through the REST API using `cloudtemplateIpsecTunnelSourceInterface`.

The following example shows a situation where only interface 3 is used as the originating interface:

```
<cloudtemplateIpSecTunnel peeraddr="173.36.19.2" preSharedKey="def" poolname="pool1">
  <cloudtemplateIpSecTunnelSourceInterface sourceInterfaceId="3" />
</cloudtemplateIpSecTunnel>
```

The following example shows a situation where both interfaces 2 and 3 are used as the originating interfaces:

```
<cloudtemplateIpSecTunnel peeraddr="173.36.19.2" preSharedKey="def" poolname="pool1">
  <cloudtemplateIpSecTunnelSourceInterface sourceInterfaceId="2" />
  <cloudtemplateIpSecTunnelSourceInterface sourceInterfaceId="3" />
</cloudtemplateIpSecTunnel>
```

Guidelines and Limitations

- Increasing the number of interfaces increases the demand of tunnel inner local IP addresses.
- The IPsec tunnel source interfaces feature is supported only with the IKEv2 configuration.

General Guidelines and Limitations for Cisco Cloud APIC

This section contains the guidelines and limitations for Cisco Cloud APIC.

- Inter-site (VRF-to-VRF) traffic is not supported if one of the VRFs is present as an attachment in a different VRF group (hub network). For example, consider the following scenario:
 - VRF-1 is stretched across different sites (Azure and AWS). In the AWS site, VRF-1 is in VRF group 1.

- VRF-2 is present in a different VRF group (VRF group 2).

In this scenario, traffic from VRF-2 to VRF-1 across sites is not supported, since the contracts between the VRFs will be implicitly allowing traffic between different VRF groups as well. Traffic across different VRF groups (hub networks) is not supported.

- You cannot stretch more than one VRF between on-prem and the cloud while using inter-VRF route leaking in the CCRs (cloud routers). For example, in a situation where VRF1 with EPG1 is stretched and VRF2 with EPG2 is also stretched, EPG1 cannot have a contract with EPG2. However, you can have multiple VRFs in the cloud, sharing one or more contracts with one on-premises VRF.
- Set the BD subnet for on-premises sites as advertised externally to advertise to the CSR1kv on the cloud.
- The default AWS security group (SG) rules limit only permits 2 CCRs per region and only 2 regions can deploy CCRs (a total maximum of 4 CCRs). To deploy more CCRs, increase the AWS SG rule limit to 120 or more. We recommend increasing the rule limit to 500.
- When configuring an object for a tenant, first check for any stale cloud resources in AWS. A stale configuration might be present if it was not cleaned properly from the previous Cisco Cloud APIC instances that managed the account.



Note It takes some time for Cisco Cloud APIC to detect the stale cloud resources after adding the tenant account ID.

To check for and clean up stale cloud resources:

1. Click the **Navigation menu** > **Application Management** > **Tenants**. The **Tenants** summary table appears in the work pane with a list of tenants as rows in a summary table.
2. Double click the tenant you are creating objects for. The **Overview**, **Cloud Resources**, **Application Management**, **Statistics**, and **Event Analytics** tabs appear.
3. Click the **Cloud Resources** > **Actions** > **View Stale Cloud Objects**. The **Stale Cloud Objects** dialog box appears.
4. If you see any stale objects, click to place a check mark in the **Automatically Clean Up Stale Cloud Objects** check box.
5. Click **Save**. The Cisco Cloud APIC automatically cleans up stale cloud objects.



Note To disable the automatic cleanup, follow steps 1 - 4 and click the **Automatically Clean Up Stale Cloud Objects** check box to remove the check mark.

- Cisco Cloud APIC tries to manage the AWS resources that it created. It does not attempt to manage resources created by other applications, other than listing existing resources as inventory. At the same time, it is also expected that AWS IAM users in the AWS infra tenant account, and the other tenant accounts, do not disturb the resources that Cisco Cloud APIC creates. For this purpose, all resources Cisco Cloud APIC creates on AWS has at least one of these two tags:

- AciDnTag

- AciOwnerTag

Cisco Cloud APIC must prevent AWS IAM users who have access to create, delete, or update EC2, or any other resources, from accessing or modifying the resources that Cisco Cloud APIC created and manages. Such restrictions should apply on both the infra tenant and other user tenant accounts. AWS account administrators should utilize the above two tags to prevent their unintentional access and modifications. For example, you can have an access policy like the following to prevent access to resources managed by Cloud APIC:

```
{
  "Effect": "Deny",
  "Action": [
    "ec2:*"
  ],
  "Resource": "*",
  "Condition": {
    "StringLike": {"ec2:ResourceTag/AciDnTag": "*"}
  }
}
```

- When configuring shared L3Out:
 - An on-premises L3Out and cloud EPGs cannot be in tenant common.
 - If an on-premises L3Out and a cloud EPG are in different tenants, define a contract in tenant common. The contract cannot be in the on-premises site or the cloud tenant.
 - Specify the CIDR for the cloud EPG in the on-premises L3Out external EPGs (l3extInstP).
 - When an on-premises L3Out has a contract with a cloud EPG in a different VRF, the VRF in which the cloud EPG resides cannot be stretched to the on-premises site and cannot have a contract with any other VRF in the on-premises site.
 - When configuring an external subnet in an on-premises external EPG:
 - Specify the external subnet as a non-zero subnet.
 - The external subnet cannot overlap with another external subnet.
 - Mark the external subnet with a shared route-control flag to have a contract with a cloud EPG.
 - The external subnet that is marked in the on-premises external EPG should have been learned through the routing protocol in the L3Out or created as a static route.
- When mapping availability zones, choose only a or b in Cisco Cloud APIC. Internally, the zone-mapping function maps this to actual availability zones in AWS.



Note The mapping works in alphabetical order. The availability zones are sorted alphabetically and then the function picks the first two and associates them to a zone a and b in Cisco Cloud APIC.

- Configuring ASN 64512 for cloud routers causes BGP sessions to not work between cloud routers and AWS virtual private gateways.
- For the total supported scale, see the following *Scale Supported* table:



-
- Note** With the scale that is specified in the *Scale Supported* table:
- You can have only 4 total managed regions.
 - You can have only CCRs in 2 regions, 2 * 2 CCRs. This is irrespective of AWS SG rule limit.
-

Table 5: Scale Supported

Component	Number Supported
Tenants	20
Applications	500
EPGs	500
Cloud Endpoints	1000
VRFs	20
Cloud Context Profiles	40
Contracts	1000
Service Graphs	200
Service Devices	100

About the Cisco Cloud APIC GUI

The Cisco Cloud APIC GUI is categorized into groups of related windows. Each window enables you to access and manage a particular component. You move between the windows using the **Navigation** menu that is located on the left side of the GUI. When you hover your mouse over any part of the menu, the following list of tab names appear: **Dashboard**, **Topology**, **Application Management**, **Cloud Resources**, **Operations**, **Infrastructure**, and **Administrative**.

Each tab contains a different list of subtabs, and each subtab provides access to a different component-specific window. For example, to view the tenant-specific window, hover your mouse over the **Navigation** menu and click **Application Management > Tenants**. From there, you can use the **Navigation** menu to view the details of another component. For example, you can navigate to the **Availability Zones** window from **Tenants** by clicking **Cloud Resources > Availability Zones**.

The **Intent** menu bar icon enables you to create a component from anywhere in the GUI. For example, to create a tenant while viewing the **Availability Zones** window, click the **Intent** icon. A dialog appears with a search box and a drop-down list. When you click the drop-down list and choose **Application Management**, a list of options, including the **Tenant** option, appears. When you click the **Tenant** option, the **Create Tenant** dialog appears displaying a group of fields that are required for creating the tenant.

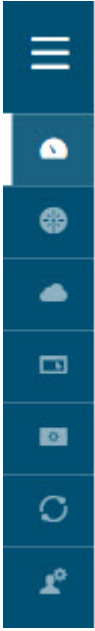
For more information about the GUI icons, see [Understanding the Cisco Cloud APIC GUI Icons, on page 15](#)

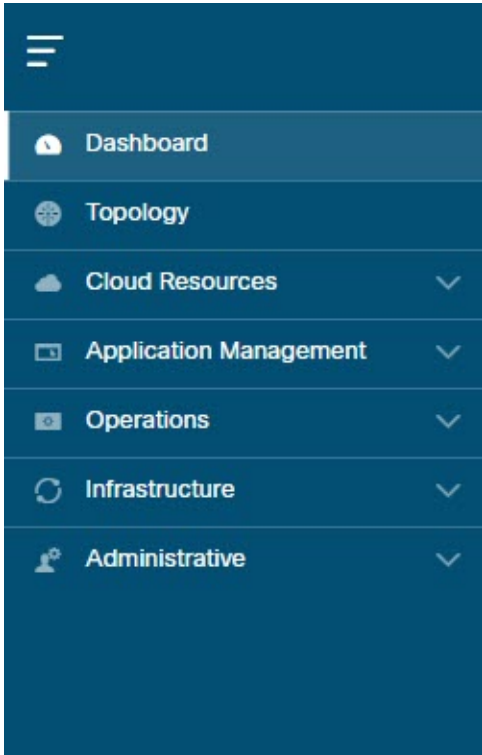


For more information about configuring Cisco Cloud APIC components, see [Configuring Cisco Cloud APIC Components, on page 45](#)







Understanding the Cisco Cloud APIC GUI Icons

This section provides a brief overview of the commonly used icons in the Cisco Cloud APIC GUI.

Table 6: Cisco Cloud APIC GUI Icons

Icon	Description
<p><i>Figure 1: Navigation Pane (Collapsed)</i></p> 	<p>The left side of the GUI contains the Navigation pane, which collapses and expands. To expand the pane, hover your mouse icon over it or click the menu icon at the top. When you click the menu icon, the Navigation pane locks in the open position. To collapse it, click the menu icon again. When you expand the Navigation pane by hovering the mouse icon over the menu icon, you collapse the Navigation pane by moving the mouse icon away from it.</p> <p>When expanded, the Navigation pane displays a list of tabs. When clicked, each tab displays a set of subtabs that enable you to navigate between the Cisco Cloud APIC component windows.</p>

Icon	Description
<p data-bbox="100 289 428 315">Figure 2: Navigation Pane (Expanded)</p> 	<p data-bbox="771 289 1466 352">The Cisco Cloud APIC component windows are organized in the Navigation pane as follows:</p> <ul data-bbox="808 369 1481 1094" style="list-style-type: none"> • Dashboard Tab—Displays summary information about the Cisco Cloud APIC components. • Topology Tab—Displays a topographical map of managed regions. • Cloud Resources Tab—Displays information about regions, availability zones, VPCs, routers, security groups, endpoints, instances, and cloud services (and target groups). • Application Management Tab—Displays information about tenants, application profiles, EPGs, contracts, filters, VRFs, service graphs, devices, and cloud context profiles. • Operations Tab—Displays information about event analytics, active sessions, backup & restore policies, tech support policies, firmware management, schedulers, and remote locations. • Infrastructure Tab—Displays information about the system configuration, inter-region connectivity, and on-premises connectivity. • Administrative Tab—Displays information about authentication, event analytics, security, local and remote users, and smart licensing. <p data-bbox="771 1125 1466 1188">Note For more information about the contents of these tabs, see Viewing System Details, on page 123</p>
<p data-bbox="100 1228 380 1253">Figure 3: Search Menu-Bar Icon</p> 	<p data-bbox="771 1228 1481 1318">The search menu-bar icon displays the search field, which enables you to search for any object by name or any other distinctive fields.</p>
<p data-bbox="100 1432 367 1457">Figure 4: Intent Menu-Bar Icon</p> 	<p data-bbox="771 1432 1466 1495">The Intent icon appears in the menu bar between the search and the help icons.</p> <p data-bbox="771 1514 1481 1671">When clicked, the Intent dialog appears (see below). The Intent dialog enables you to create a component from any window in the Cisco Cloud APIC GUI. When you create or view a component, a dialog box opens and hides the Intent icon. Close the dialog box to access the Intent icon again.</p> <p data-bbox="771 1690 1481 1753">For more information about creating a component, see Configuring Cisco Cloud APIC Components, on page 45.</p>

Icon	Description
<p>Figure 5: Intent (What do you want to do?) Dialog Box</p> 	<p>The Intent (What do you want to do?) dialog box contains a search box and a drop-down list. The drop-down list enables you to apply a filter for displaying specific options. The search box enables you to enter text for searching through the filtered list.</p>
<p>Figure 6: Feedback Icon</p> 	<p>The feedback icon appears in the menu bar between the Intent and the bookmark icons.</p> <p>When clicked, the feedback panel appears.</p>
<p>Figure 7: Bookmark Icon</p> 	<p>The bookmark icon appears in the menu bar between the feedback and the system tools icons.</p> <p>When clicked, the current page is bookmarked on your system.</p>
<p>Figure 8: System Tools Menu-Bar Icon</p> 	<p>The system tools menu-bar icon provides the following options:</p> <ul style="list-style-type: none"> • About—Display the Cisco Cloud APIC version. • ObjectStore Browser—Open the Managed Object Browser, or Visore, which is a utility that is built into Cisco Cloud APIC that provides a graphical view of the managed objects (MOs) using a browser.
<p>Figure 9: Help Menu-Bar Icon</p> 	<p>The help menu-bar icon shows the About Cloud APIC menu option, which provides the version information for the Cloud APIC. The help menu-bar icon also shows the Help Center and Welcome Screen menu options.</p>
<p>Figure 10: User Profile Menu-Bar Icon</p> 	<p>The user profile menu-bar icon provides the following options:</p> <ul style="list-style-type: none"> • User Preferences—Allows you to set the time format (Local or UTC) and enable or disable the Welcome Screen at login. • Change Password—Enables you to change the password. • Change SSH Key—Enables you to change the SSH key. • Change User Certificate—Enables you to change the user certificate. • Logout—Enables you to log out of the GUI.



CHAPTER 3

Cisco Cloud APIC Policy Model

- [About the ACI Policy Model, on page 19](#)
- [Policy Model Key Characteristics, on page 19](#)
- [Logical Constructs, on page 20](#)
- [The Cisco ACI Policy Management Information Model, on page 21](#)
- [Tenants, on page 23](#)
- [Cloud Context Profile, on page 24](#)
- [VRFs, on page 34](#)
- [Cloud Application Profiles, on page 35](#)
- [Cloud Endpoint Groups, on page 35](#)
- [Contracts, on page 37](#)
- [About the Cloud Template, on page 39](#)
- [Managed Object Relations and Policy Resolution, on page 42](#)
- [Default Policies, on page 43](#)
- [Shared Services, on page 44](#)

About the ACI Policy Model

The ACI policy model enables the specification of application requirements policies. The Cisco Cloud APIC automatically renders policies in the cloud infrastructure. When you or a process initiates an administrative change to an object in the cloud infrastructure, the Cisco Cloud APIC first applies that change to the policy model. This policy model change then triggers a change to the actual managed item. This approach is called a model-driven framework.

Policy Model Key Characteristics

Key characteristics of the policy model include the following:

- As a model-driven architecture, the software maintains a complete representation of the administrative and operational state of the system (the model). The model applies uniformly to cloud infrastructure, cloud infrastructure, services, system behaviors, and virtual devices attached to the network.
- The logical and concrete domains are separated; the logical configurations are rendered into concrete configurations by applying the policies in relation to the available resources. No configuration is carried

out against concrete entities. Concrete entities are configured implicitly as a side effect of the changes to the Cisco Cloud policy model.


- The system prohibits communications with newly connected endpoints until the policy model is updated to include the new endpoint.
- Network administrators do not configure logical system resources directly. Instead, they define logical (hardware-independent) configurations and the Cisco Cloud APIC policies that control different aspects of the system behavior.

Managed object manipulation in the model relieves engineers from the task of administering isolated, individual component configurations. These characteristics enable automation and flexible workload provisioning that can locate any workload anywhere in the infrastructure. Network-attached services can be easily deployed, and the Cisco Cloud APIC provides an automation framework to manage the lifecycle of those network-attached services.

Logical Constructs

The policy model manages the entire cloud infrastructure, including the infrastructure, authentication, security, services, applications, cloud infrastructure, and diagnostics. Logical constructs in the policy model define how the cloud infrastructure meets the needs of any of the functions of the cloud infrastructure. The following figure provides an overview of the ACI policy model logical constructs.

Figure 11: ACI Policy Model Logical Constructs Overview




cloud infrastructure-wide or tenant administrators create predefined policies that contain application or shared resource requirements. These policies automate the provisioning of applications, network-attached services, security policies, and tenant subnets, which puts administrators in the position of approaching the resource pool in terms of applications rather than infrastructure building blocks. The application needs to drive the networking behavior, not the other way around.

The Cisco ACI Policy Management Information Model

The cloud infrastructure comprises the logical components as recorded in the Management Information Model (MIM), which can be represented in a hierarchical management information tree (MIT). The Cisco Cloud APIC runs processes that store and manage the information model. Similar to the OSI Common Management Information Protocol (CMIP) and other X.500 variants, the Cisco Cloud APIC enables the control of managed resources by presenting their manageable characteristics as object properties that can be inherited according to the location of the object within the hierarchical structure of the MIT.

Each node in the tree represents a managed object (MO) or group of objects. MOs are abstractions of cloud infrastructure resources. An MO can represent a concrete object, such as a cloud router, adapter, or a logical object, such as an application profile, cloud endpoint group, or fault. The following figure provides an overview of the MIT.

Figure 12: Cisco ACI Policy Management Information Model Overview



The hierarchical structure starts with the policy universe at the top (Root) and contains parent and child nodes. Each node in the tree is an MO and each object in the cloud infrastructure has a unique distinguished name (DN) that describes the object and locates its place in the tree.

The following managed objects contain the policies that govern the operation of the system:

- A tenant is a container for policies that enable an administrator to exercise role-based access control. The system provides the following four kinds of tenants:
 - The administrator defines user tenants according to the needs of users. They contain policies that govern the operation of resources such as applications, databases, web servers, network-attached storage, virtual machines, and so on.
 - Although the system provides the common tenant, it can be configured by the cloud infrastructure administrator. It contains policies that govern the operation of resources accessible to all tenants, such as firewalls, load balancers, Layer 4 to Layer 7 services, intrusion detection appliances, and so on.



Note As of the Cisco Application Policy Infrastructure Controller (APIC) Release 4.1(1), the Cisco Cloud APIC only supports load balancers as a Layer 4 to Layer 7 service.

- The infrastructure tenant is provided by the system but can be configured by the cloud infrastructure administrator. It contains policies that govern the operation of infrastructure resources. It also enables a cloud infrastructure provider to selectively deploy resources to one or more user tenants. Infrastructure tenant policies are configurable by the cloud infrastructure administrator.
- The cloud infra policies enable you to manage on-premises and inter-region connectivity when setting up the Cisco Cloud APIC. For more information, see the *Cisco Cloud APIC Installation Guide*.
- Cloud inventory is a service that enables you to view different aspects of the system using the GUI. For example, you can view the regions that are deployed from the aspect of an application or the applications that are deployed from the aspect of a region. You can use this information for cloud resource planning and troubleshooting.
- Layer 4 to Layer 7 service integration lifecycle automation framework enables the system to dynamically respond when a service comes online or goes offline. For more information, see [Deploying Layer 4 to Layer 7 Services, on page 135](#)


- Access, authentication, and accounting (AAA) policies govern user privileges, roles, and security domains of the Cisco Cloud ACI cloud infrastructure. For more information, see [Cisco Cloud APIC Security](#), on page 163

The hierarchical policy model fits well with the REST API interface. When invoked, the API reads from or writes to objects in the MIT. URLs map directly into distinguished names that identify objects in the MIT. Any data in the MIT can be described as a self-contained structured tree text document encoded in XML or JSON.

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 13: Tenants



Tenants can be isolated from one another or can share resources. The primary elements that the tenant contains are filters, contracts, Virtual Routing and Forwarding (VRF) instances, cloud context profiles, AWS provider configurations, and cloud application profiles that contain cloud endpoint groups (cloud EPGs). Entities in the tenant inherit its policies. VRFs are also known as contexts; each VRF can be associated with multiple cloud context profiles. A cloud context profile in conjunction with a VRF and a region represents the AWS VPC in that region.

Tenants are logical containers for application policies. The cloud infrastructure can contain multiple tenants. You must configure a tenant before you can deploy any Layer 4 to Layer 7 services. The ACI cloud infrastructure supports IPv4 and dual-stack configurations for tenant networking.

Cloud Context Profile

The cloud context profile contains information on the following Cisco Cloud APIC components:

- Availability zones and regions
- CIDRs
- CCRs
- Endpoints
- EPGs
- Virtual Networks
- VRFs

The following sections provide additional information on some of the components that are part of the cloud context profile.

CCR

The CCR is a virtual router that delivers comprehensive WAN gateway and network services into virtual and cloud environments. The CCR enables enterprises to extend their WANs into provider-hosted clouds. Two CCRs are required for Cisco Cloud APIC solution.

The type of CCR used with the Cisco Cloud APIC varies depending on the release:

- For releases prior to release 25.0(3), the **Cisco Cloud Services Router 1000v** is used with the Cisco Cloud APIC. For more information on this type of CSR, see the [Cisco Cloud Services Router 1000v documentation](#).
- For release 25.0(3) and later, the **Cisco Catalyst 8000V** is used with the Cisco Cloud APIC. For more information on this type of CCR, see the [Cisco Catalyst 8000V Edge software documentation](#).

About the Cisco Catalyst 8000V

Beginning with release 25.0(3), Cisco Cloud APIC moves from the Cisco Cloud Services Router 1000v to the Cisco Catalyst 8000V. Following are updates that are specific to the Cisco Catalyst 8000V.

- [Licensing, on page 24](#)
- [Throughput, on page 25](#)

Licensing

Beginning with release 25.0(4), the Cisco Catalyst 8000V on Cisco Cloud APIC supports the following licensing models:

1. **Bring Your Own License (BYOL)** Licensing Model
2. **Pay As You Go (PAYG)** Licensing Model



Note For releases prior to 25.0(4), the Cisco Catalyst 8000V on Cisco Cloud APIC supports only the **Bring Your Own License (BYOL)** licensing model.

BYOL Licensing Model

The BYOL licensing model on Cisco Catalyst 8000V which requires you to purchase your Catalyst 8000V Cisco DNA license from Cisco and deploy it in the cloud.

- For instructions on subscribing to one of the tier-based Cisco Catalyst 8000V licenses, see [Cisco Catalyst 8000V Edge Software](#).
- For more information on different throughputs based on the tiers, see [Throughput, on page 25](#).

Cisco Cloud APIC makes use of the “Cisco DNA Advantage” subscription. For features supported by the “Cisco DNA Advantage” subscription, see [Cisco DNA Software SD-WAN and Routing Matrices](#).

PAYG Licensing Model

Beginning with the 25.0(4) release, Cisco Cloud APIC supports Pay-As-You-Go (PAYG) Licensing Model on Cisco Catalyst 8000V which allows users to deploy a Catalyst 8000V instance in the cloud based on the VM size and purchase the usage on an hourly basis.

As you completely depend on the VM size to get the throughput, the PAYG licensing model can be enabled only by first un-deploying the current Cisco Catalyst 8000V and then re-deploying it using the First Time Set Up with the new VM size. For more information, see the chapter "Configuring Cisco Cloud APIC Using the Setup Wizard" in the [Cisco Cloud APIC for AWS Installation Guide](#).



Note The procedure for switching between licenses can also be used if you would like to switch between the two licensing types available.



Note There are two PAYG options for consuming licenses in the AWS marketplace: **Catalyst 8000V Cisco DNA Essentials** and **Catalyst 8000V Cisco DNA Advantage**. Cisco Cloud APIC will make use of **Catalyst 8000V Cisco DNA Advantage**. For features supported by the “Cisco DNA Advantage” subscription, see [Cisco DNA Software SD-WAN and Routing Matrices](#)

Throughput

Beginning with release 25.0(4), the Cisco Catalyst 8000V on Cisco Cloud APIC supports the following licensing models:

1. **Bring Your Own License (BYOL)** Licensing Model
2. **Pay As You Go (PAYG)** Licensing Model



Note For releases prior to 25.0(4), the Cisco Catalyst 8000V on Cisco Cloud APIC supports only the **Bring Your Own License (BYOL)** licensing model.

1. Bring Your Own License (BYOL)

For this model, the Cisco Catalyst 8000V supports tier-based (T0/T1/T2/T3) throughput options. The following table lists what AWS EC2 instance is used for different router throughput settings for the Cisco Catalyst 8000V:

CCR Throughput	AWS EC2 Instance
T0 (up to 15M throughput)	c5.xlarge
T1 (up to 100M throughput)	c5.xlarge
T2 (up to 1G throughput)	c5.xlarge
T3 (up to 10G throughput)	c5.9xlarge

Tier2 (T2) is the default throughput supported by Cisco Cloud APIC.

The following table shows the mapping of throughput from the older Cisco Cloud Services Router 1000v routers to the newer Cisco Catalyst 8000V routers:

Throughput on Cisco Cloud Services Router 1000v	Throughput on Cisco Catalyst 8000V
10M	T0 (up to 15M throughput)
50M	T1 (up to 100M throughput)
100M	T1 (up to 100M throughput)
250M	T2 (up to 1G throughput)
500M	T2 (up to 1G throughput)
1G	T2 (up to 1G throughput)
2.5G	T3 (up to 10G throughput)
5G	T3 (up to 10G throughput)
7.5G	T3 (up to 10G throughput)
10G	T3 (up to 10G throughput)

2. Pay-As-You-Go Licensing Model

For this model, Cisco Cloud APIC supports a range of AWS EC2 instances for cloud networking needs powered by Cisco's Catalyst 8000V virtual router.

The table below shows the cloud instance type supported by Cisco Cloud APIC on AWS.

AWS EC2 Instance	CCR Throughput	vCPUs	Memory
c5.xlarge	up to 5 Gigabit throughput	4	8 GiB
c5.2xlarge	up to 10 Gigabit throughput	8	16 GiB

AWS EC2 Instance	CCR Throughput	vCPUs	Memory
c5.4xlarge	up to 10 Gigabit throughput	16	32 GiB
c5.9xlarge	up to 10 Gigabit throughput	36	72 GiB
c5n.xlarge	up to 25 Gigabit throughput	4	10.5 GiB
c5n.2xlarge	up to 25 Gigabit throughput	8	21 GiB
c5n.4xlarge	up to 25 Gigabit throughput	16	42 GiB
c5n.9xlarge	up to 50 Gigabit throughput	36	96 GiB

Private IP Address Support for Cisco Cloud APIC and CCR in AWS



Note For Azure, support for private IP addresses for Cisco Cloud APIC and CCRs became available in release 5.1(2). For AWS, this support is available beginning with release 5.2(1).

For AWS, prior to release 5.2(1), Cisco Cloud Router (CCR) interfaces were assigned both public and private IP address by Cisco Cloud APIC. Beginning with release 5.2(1), CCR interfaces are assigned private IP addresses only and assignment of public IP addresses to CCR interfaces is optional. Private IP addresses are always assigned to all the interfaces of a CCR. The private IP address of GigabitEthernet1 of a CCR is used as BGP and OSPF router IDs.

To enable CCR private IP addresses for inter-site connectivity, where you are disabling public IP addresses for the CCR interfaces, see the [Managing Regions \(Configuring a Cloud Template\) Using the Cisco Cloud APIC GUI, on page 110](#) procedure.

For AWS, prior to release 5.2(1), the management interface of the Cisco Cloud APIC was assigned a public IP address and a private IP address. Beginning with release 5.2(1), a private IP address is assigned to the management interface of the Cisco Cloud APIC and assigning a public IP address is optional. To disable public IP to the Cisco Cloud APIC so that a private IP address is used for connectivity, see the *Deploying the Cloud APIC in AWS* procedure in the *Cisco Cloud APIC for AWS Installation Guide*, Release 5.2(1) or later.

Restrictions for CCR with Private IP Address

When public IP is disabled, the underlay inter-site connectivity cannot be Public internet because Public Internet requires a public IP address. The underlay inter-site connectivity can only be one of the following:

- Private connection for connecting to an on-premise ACI site, which is through AWS Direct Connect or Azure Express Route
- Cloud backbone for connecting to a Cisco Cloud APIC site of the same cloud provider, which is through AWS VPC Peering or Azure Vnet Peering

Communicating to External Sites From Regions Without a CCR

Prior to release 5.2(1), for traffic to pass through to an external site, the region where the traffic is passing through must have a CCR deployed. The CCR advertises the CIDRs that are local to that region. If an EPG in a region has a contract with an external site, then that region must have a CCR deployed in order to communicate with that external site.


Beginning with release 5.2(1), you can now have communication with an external site from regions without a CCR. This is accomplished by making use of the AWS Transit Gateway feature, which became available for Cisco Cloud APIC in release 5.0(1). When you enable the AWS Transit Gateway feature on Cisco Cloud APIC, you also enable peering between all managed regions on AWS. In this way, the AWS Transit Gateway peering feature allows the Cisco Cloud APIC to address the issue of communicating with external sites from regions without a CCR. It addresses this issue by having traffic rerouted to a region with a CCR.

Using the AWS Transit Gateway feature, when traffic from a region without a CCR tries to egress out of a site, this traffic will be routed to the infra VPC for the closest region with a CCR. After the traffic is rerouted to that region's infra VPC, that CCR is used to egress out the packet. For ingress traffic, packets coming from an external site can reach any region's CCR and then be routed to the destination region using the AWS Transit Gateway peering in the ingress data path.

In these situations, traffic is rerouted automatically when the system recognizes that external traffic is egressing or ingressing through a region without a CCR. However, you must have the following components configured in order for the system to automatically perform this rerouting task:

- AWS Transit Gateway must be configured. If AWS Transit Gateway is not already configured, see [Increasing Bandwidth Between VPCs by Using AWS Transit Gateway](#) for those instructions.
- CCRs must be deployed in at least one region. Even though this enhancement allows you to communicate with an external site from a region that *does not* contain a CCR, in order to do this, you must have another separate region that *does* contain a CCR so that traffic can be rerouted from the region without a CCR to the region with a CCR.

The following figure shows an example scenario where traffic is rerouted automatically when the system recognizes that external traffic is egressing from a region without a CCR.



503711

The following occurs when the Cisco Cloud APIC recognizes that Region 2 does not have a CCR, but traffic is egressing out to an external site (shown with the green arrow and circles):

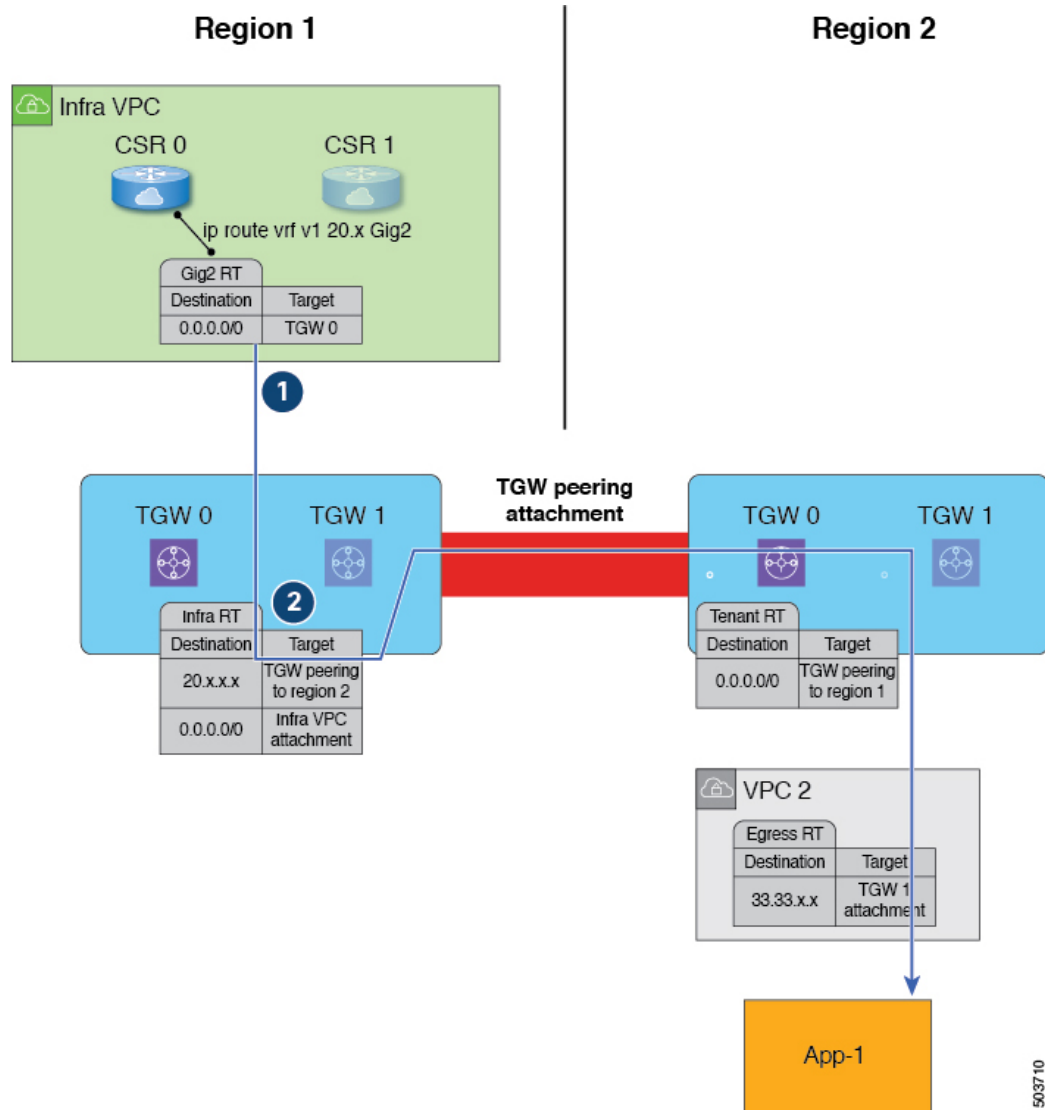
1. Traffic flow begins egressing out from the CIDR in App-1 in Region 2 to a remote site. Note that the endpoint is configured with the appropriate outbound rule to allow the remote site CIDR.
2. The VPC 2 egress route table has the remote site CIDR, which then has the AWS Transit Gateway as the next hop. The AWS Transit Gateway information is programmed automatically based on the configured contracts.
3. A 0.0.0.0/0 route is inserted in the AWS Transit Gateway route table, which essentially tells the system to take the AWS Transit Gateway peering attachment if traffic is egressing out to an external site but there is no CCR in this region. In this situation, the AWS Transit Gateway peering attachment goes to another region that does have a CCR (Region 1 in the example scenario). The region with a CCR that will be used is chosen based on geographical proximity to the region that does not have a CCR.

4. The packet is first forwarded to the infra VPC in the region with the CCR (Region 1), and is then forwarded to the gigabit ethernet network interface on the CCR that is associated with the appropriate VRF group.
5. The gigabit 2 inbound security group on the CCR in Region 1 is configured to allow traffic from the remote region VPC.

It's useful to note that in the egress example shown above:

- For steps 1 and 2, there is no change from pre-release 5.2(1) behavior.
- Step 3 is behavior that is new and unique to this feature in release 5.2(1), where the redirect occurs to the TGW peering attachment from the region without a CCR to the region with a CCR. In addition, step 3 occurs on the AWS cloud.
- Steps 4 and 5 would normally occur in Region 2 before release 5.2(1), but only if Region 2 had a CCR. However, because Region 2 does not have a CCR, with this feature in release 5.2(1), these steps are taking place in Region 1 where a CCR is present.

The following figure shows an example scenario where traffic is rerouted automatically when the system recognizes that external traffic is ingressing to a region without a CCR.



The following occurs when the Cisco Cloud APIC recognizes that Region 2 does not have a CCR, but traffic is ingressing in from an external site to a CIDR in App-1 in Region 2 (shown with the blue arrow and circles):

1. Normally, the CCR in Region 1 would only advertise the CIDRs that are local to that region. However, with this enhancement that is part of release 5.2(1), all CCRs in all regions now advertise CIDRs from all remote regions. Therefore, in this example, the CCR in Region 1 will also advertise the CIDRs that are in Region 2 (assuming AWS Transit Gateway peering is configured between both regions and the contracts are configured correctly). In this situation, the traffic ingresses in from an external site to the CCR in Region 1, where the CCR in Region 1 advertises the static route for the remote region VPC CIDRs.
2. The infra route table (the AWS Transit Gateway route table in Region 1) has the next hop to the AWS Transit Gateway peering attachment to Region 2.

It's useful to note that in the ingress example shown above:

- Both steps (steps 1 and 2) in the ingress example shown above are new and unique to this feature in release 5.2(1).

- Step 1 in the ingress example shows configurations programmed on the CCR.
- Step 2 in the ingress example occurs on the AWS cloud.

Support for ECMP Forwarding from Remote Sites for CCRs

CCRs in a cloud will typically receive more than one path for a prefix. Prior to release 5.2(1), there was no support for data forwarding from CCRs using Equal Cost Multiple Path (ECMP), even though the CCR receives multiple paths.

Beginning with release 5.2(1), support is now available for ECMP with CCRs, where traffic from CCRs will be forwarded to all ECMP paths received from a destination site. This support is automatically enabled with release 5.2(1) and requires no manual configuration to enable this support.

Preference For Routes to CCRs in Regions with Local CIDRs

Every CIDR that is configured is local to a specific region. With multiple regions in a cloud, CCRs from all regions advertise the CIDRs for redundancy. Prior to release 5.2(1), CCRs from all regions advertised the CIDRs with the same preference. This can cause a remote cloud site or an on-prem site to install the path to a CIDR through a region where the CIDR is not local. This, in turn, could result in a packet taking a longer route than necessary.

Beginning with release 5.2(1), CCRs from multiple regions will continue to advertise the CIDRs, but CCRs from the region where the CIDR is local will advertise with a higher preference. This causes the on-premises site or the remote cloud site to direct traffic directly to the region where the CIDR is local. If the CCRs in the local region fail, the paths from the other regions can be used for data forwarding.

Availability Zones

Prior to release 25.0(2), Cisco Cloud APIC supports only two availability zones per region in AWS, where Cisco Cloud APIC creates two availability zones, called **virtual** availability zones, for each region using the format <region-name>a and <region-name>b. For example, under the `us-west-1` region, Cisco Cloud APIC creates the two virtual availability zones `us-west-1a` and `us-west-1b`.

Beginning with release 25.0(2), the **cloud** availability zone is now supported, which allows for multiple availability zones in each AWS region with Cisco Cloud APIC.

- To view the **virtual** availability zones for your Cisco Cloud APIC, navigate to **Cloud Resources > Availability Zones**, then click the **Virtual Availability Zones** tab.

Name	Cloud Availability Zone	Tenants	Application Management		Cloud Resources		
			App. Profiles	EPGs	VPCs	Routers	Endpoints
af-south-1-1a region-af-south-1		N/A	N/A	N/A	N/A	0	N/A
af-south-1-1b region-af-south-1		N/A	N/A	N/A	N/A	0	N/A
ap-east-1-1a region-ap-east-1		N/A	N/A	N/A	N/A	0	N/A

- To view the **cloud** availability zones for your Cisco Cloud APIC, navigate to **Cloud Resources > Availability Zones**, then click the **Cloud Availability Zones** tab.


Name	Tenants	Application Management		Cloud Resources		
		App. Profiles	EPGs	VPCs	Routers	Endpoints
us-east-1a AWS > infra > us-east-1	N/A	N/A	N/A	N/A	0	N/A
us-east-1b AWS > infra > us-east-1	N/A	N/A	N/A	N/A	0	N/A
us-east-1c AWS > infra > us-east-1	N/A	N/A	N/A	N/A	0	N/A

Migrating from Virtual Availability Zones to Cloud Availability Zones

If you have deployments that you configured prior to release 25.0(2), where you have virtual availability zones configured, when you upgrade to release 25.0(2), we recommend that you migrate from the older virtual availability zones to the newer cloud availability zones after you have upgraded to release 25.0(2).

- You can migrate individual subnets or all of the subnets in a CIDR block range as part of the availability zone migration.
- Migrating from older virtual availability zones to the newer cloud availability zones will not cause have any functional impact, such as a traffic drop, in the cloud resources in AWS.



Note The following steps describe how to migrate from virtual availability zones to cloud availability zones through the cloud context profile, but you can also migrate availability zones by clicking the Intent icon () and selecting **Availability Zone Configuration Migration**.

To migrate from virtual availability zones to cloud availability zones:

1. Navigate to the cloud context profile that was configured previously with the older virtual availability zones.

In the left nav pane, navigate to **Application Management > Cloud Context Profiles**, then locate the cloud context profile that was configured previously with the older virtual availability zones.

2. Double-click on that cloud context profile.

The details panel for that cloud context profile appears with the **Overview** tab selected automatically.

View the entries in the **Availability Zone** column in the **Overview** tab to determine if you have virtual availability zones in this cloud context profile that you can migrate to cloud availability zones.

3. Click **Actions > Migrate Subnet Configuration**.

The **Availability Zone Configuration Migration** window appears.

4. Select the subnets associated with the virtual availability zones that you want to migrate to cloud availability zones.

- All of the subnets listed in this window that are associated with virtual availability zones will be selected by default. Manually deselect any subnets associated with virtual availability zones that you do not want to migrate to cloud availability zones.

- For each virtual availability zone that will be migrated over to cloud availability zones, make a note of the entry in the Cloud Availability Zones column to determine the new availability zone value for that subnet, if necessary.

5. Click **Migrate Subnet Configuration**.

The selected virtual availability zones are migrated to cloud availability zones.

Guidelines and Limitations


Following are the guidelines and limitations for support of multiple availability zones:

- Support for cloud availability zones, where you can have more than two availability zones, is available for user tenants only. Infra tenants will continue to use virtual availability zones which have a limit of two availability zones.

VRFs

A Virtual Routing and Forwarding (VRF) object (`fvCtx`) or context is a tenant network (called a private network in the Cisco Cloud APIC GUI). A tenant can have multiple VRFs. A VRF is a unique Layer 3 forwarding and application policy domain. The following figure shows the location of VRFs in the management information tree (MIT) and their relation to other objects in the tenant.

Figure 14: VRFs

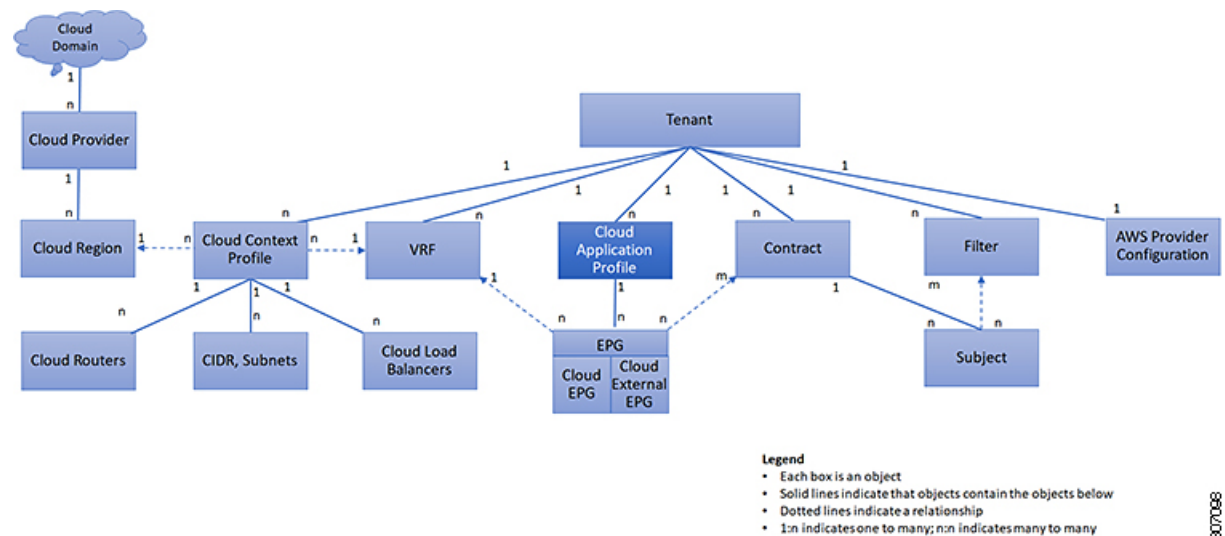


A VRF defines a Layer 3 address domain. One or more cloud context profiles are associated with a VRF. You can only associate one cloud context profile with a VRF in a given region. All the endpoints within the Layer 3 domain must have unique IP addresses because it is possible to forward packets directly between these devices if the policy allows it. A tenant can contain multiple VRFs. After an administrator creates a logical device, the administrator can create a VRF for the logical device, which provides a selection criteria policy for a device cluster. A logical device can be selected based on a contract name, a graph name, or the function node name inside the graph.

Cloud Application Profiles

A cloud application profile (c_loudAp) defines the policies, services and relationships between cloud EPGs. The following figure shows the location of cloud application profiles in the management information tree (MIT) and their relation to other objects in the tenant.

Figure 15: Cloud Application Profiles



Cloud application profiles contain one or more cloud EPGs. Modern applications contain multiple components. For example, an e-commerce application could require a web server, a database server, data located in a storage service, and access to outside resources that enable financial transactions. The cloud application profile contains as many (or as few) cloud EPGs as necessary that are logically related to providing the capabilities of an application.


Cloud EPGs can be organized according to one of the following:

- The application they provide, such as a DNS server or SAP application (see *Tenant Policy Example* in *Cisco APIC REST API Configuration Guide*).
- The function they provide (such as infrastructure)
- Where they are in the structure of the data center (such as DMZ)
- Whatever organizing principle that a cloud infrastructure or tenant administrator chooses to use

Cloud Endpoint Groups

The cloud endpoint group (cloud EPG) is the most important object in the policy model. The following figure shows where application cloud EPGs are located in the management information tree (MIT) and their relation to other objects in the tenant.

Figure 16: Cloud Endpoint Groups



A cloud EPG is a managed object that is a named logical entity that contains a collection of endpoints. Endpoints are devices that are connected to the network directly or indirectly. They have an address (identity), a location, attributes (such as version or patch level), and are virtual. Knowing the address of an endpoint also enables access to all its other identity details. Cloud EPGs are fully decoupled from the physical and logical topology. Endpoint examples include servers, virtual machines, storage services, or clients on the Internet. Endpoint membership in a cloud EPG can be dynamic or static.

The ACI cloud infrastructure can contain the following types of cloud EPGs:

- Cloud endpoint group (`cloudEPg`)
- Cloud external endpoint group (`cloudExtEPg`)

Cloud EPGs contain endpoints that have common policy requirements such as security or Layer 4 to Layer 7 services. Rather than configure and manage endpoints individually, they are placed in a cloud EPG and are managed as a group.

Policies apply to cloud EPGs, never to individual endpoints.

Regardless of how a cloud EPG is configured, cloud EPG policies are applied to the endpoints they contain.

WAN router connectivity to the cloud infrastructure is an example of a configuration that uses a static cloud EPG. To configure WAN router connectivity to the cloud infrastructure, an administrator configures a `cloudExtEPg` cloud EPG that includes any endpoints within an associated WAN subnet. The cloud infrastructure learns of the cloud EPG endpoints through a discovery process as the endpoints progress through their connectivity life cycle. Upon learning of the endpoint, the cloud infrastructure applies the `cloudExtEPg` cloud EPG policies accordingly. For example, when a WAN connected client initiates a TCP session with a server within an application (`cloudEPg`) cloud EPG, the `cloudExtEPg` cloud EPG applies its policies to that client endpoint before the communication with the (`cloudExtEPg`) cloud EPG web server begins. When the client server TCP session ends, and communication between the client and server terminates, the WAN endpoint no longer exists in the cloud infrastructure.


The Cisco Cloud APIC uses endpoint selectors to assign endpoints to Cloud EPGs. The endpoint selector is essentially a set of rules that are run against the cloud instances that are assigned to the AWS VPC managed

by Cisco ACI. Any endpoint selector rules that match endpoint instances assign that endpoint to the Cloud EPG. The endpoint selector is similar to the attribute-based microsegmentation available in Cisco ACI.

Contracts

In addition to cloud EPGs, contracts (`vzBfCP`) are key objects in the policy model. Cloud EPGs can only communicate with other cloud EPGs according to contract rules. The following figure shows the location of contracts in the management information tree (MIT) and their relation to other objects in the tenant.

Figure 17: Contracts



An administrator uses a contract to select one or more types of traffic that can pass between cloud EPGs, including the protocols and ports allowed. If there is no contract, inter-EPG communication is disabled by default. There is no contract required for intra-EPG communication; intra-EPG communication is always implicitly allowed.

Contracts govern the following types of cloud EPG communications:

- Between cloud EPGs (`cloudEPg`), both intra-tenant and inter-tenant



Note In the case of a shared service mode, a contract is required for inter-tenant communication. A contract is used to specify static routes across VRFs, although the tenant VRF does not enforce a policy.

- Between cloud EPGs and cloud external EPGs (`cloudExtEPg`)

Contracts govern the communication between cloud EPGs that are labeled providers, consumers, or both. The relationship between a cloud EPG and a contract can be either a provider or consumer. When a cloud EPG provides a contract, communication with that cloud EPG can be initiated from other cloud EPGs as long as the communication complies with the provided contract. When a cloud EPG consumes a contract, the cloud endpoints in the consuming cloud EPG may initiate communication with any cloud endpoint in a cloud EPG that is providing that contract.




Note A cloud EPG can both provide and consume the same contract. A cloud EPG can also provide and consume multiple contracts simultaneously.

Filters and Subjects Govern Cloud EPG Communications

Subject and filter managed-objects enable mixing and matching among cloud EPGs and contracts so as to satisfy various applications or service delivery requirements. The following figure shows the location of application subjects and filters in the management information tree (MIT) and their relation to other objects in the tenant.

Figure 18: Subjects and Filters



Contracts can contain multiple communication rules and multiple cloud EPGs can both consume and provide multiple contracts. A policy designer can compactly represent complex communication policies and re-use these policies across multiple instances of an application.



Note Subjects are hidden in Cisco Cloud APIC and not configurable. For rules installed in AWS, source port provided in the filter entry is not taken into account.

Subjects and filters define cloud EPG communications according to the following options:

- Filters are Layer 2 to Layer 4 fields, TCP/IP header fields such as Layer 3 protocol type, Layer 4 ports, and so forth. According to its related contract, a cloud EPG provider dictates the protocols and ports in both the in and out directions. Contract subjects contain associations to the filters (and their directions) that are applied between cloud EPGs that produce and consume the contract.



Note When a contract filter match type is `ALL`, best practice is to use the VRF unenforced mode. Under certain circumstances, failure to follow these guidelines results in the contract not allowing traffic among cloud EPGs in the VRF.

- Subjects are contained in contracts. One or more subjects within a contract use filters to specify the type of traffic that can be communicated and how it occurs. For example, for HTTPS messages, the subject specifies the direction and the filters that specify the IP address type (for example, IPv4), the HTTP protocol, and the ports allowed. Subjects determine if filters are unidirectional or bidirectional. A unidirectional filter is used in one direction. Unidirectional filters define in or out communications but not the same for both. Bidirectional filters are the same for both; they define both in and out communications.



Note For rules that are installed in AWS, the source port provided in the filter entry is not taken into account.

- ACI contracts rendered in AWS constructs are always stateful, allowing return traffic.

About the Cloud Template

The cloud template provides a template that configures and manages the Cisco Cloud APIC infra network. The template requires only the most essential elements for the configuration. From these elements, the cloud template generates a detailed configuration necessary for setting up the Cisco Cloud APIC infra network. However, it is not a one-time configuration generation—it is possible to add, modify, or remove elements of the template input. The cloud template updates the resulting configuration accordingly.

One of the central things in the AWS network configuration is the Virtual Private Cloud (VPC). AWS supports many regions worldwide and one VPC is specific to one region.

The cloud template accepts one or more region names and generates the entire configuration for the infra VPCs in those regions. They are the infra VPCs. The Cisco Cloud APIC-managed object (MO) corresponding to the AWS VPC is `cloudCtxProfile`. For every region specified in the cloud template, it generates the `cloudCtxProfile` configuration. A `cloudCtxProfile` is the topmost MO for all the configuration corresponding to a region. Underneath, it has many of other MOs organized as a tree to capture a specific configuration. A `cloudCtxProfile` MO generated by the cloud template carries `ctxProfileOwner == SYSTEM`. For the non-infra network, it is possible to configure `cloudCtxProfile` directly; in this case, `cloudCtxProfile` carries `ctxProfileOwner == USER`.

A primary property of an AWS VPC is the CIDR. Every region needs a unique CIDR. In Cisco Cloud APIC, you can provide the CIDRs for the infra VPCs. The CIDRs for the first two regions come from the Cloud Formation Template (CFT) that deploys the Cisco Cloud APIC AMI on the AWS. The `cloudApicSubnetPool` MO provides CIDRs for the additional regions directly to the Cisco Cloud APIC. In the Cisco Cloud APIC configuration, the `cloudCidr` MO, which is a child of `cloudCtxProfile`, models the CIDR.

The cloud template generates and manages a huge number of MOs in the `cloudCtxProfile` subtree including, but not limited to, the following:

- Subnets
- Association of subnets to AWS availability zones
- Cloud routers
- IP address allocation for the cloud router interfaces
- IP address allocation and configuration for tunnels
- IP address allocation and configuration for loopbacks

Without the cloud template, you would be responsible for configuring and managing these.

The *Cisco Cloud Template MO* table contains a brief summary of the inputs (MOs) to the cloud template.

Table 7: Cloud Template MOs

MO	Purpose
<code>cloudtemplateInfraNetwork</code>	The root of the cloud template configuration. Attributes include: <code>numRoutersPerRegion</code> —The number of cloud routers for each <code>cloudRegionName</code> specified under <code>cloudtemplateIntNetwork</code> .
<code>cloudtemplateProfile</code>	Configuration profile for all the cloud routers. Attributes include: <ul style="list-style-type: none"> • <code>routerUsername</code> • <code>routerPassword</code> • <code>routerThroughput</code> • <code>routerLicenseToken</code> • <code>routeDataInterfacePublicIP</code> • <code>routerMgmtInterfacePublicIP</code>
<code>cloudtemplateIntNetwork</code>	Contains a list of regions, which specify where you deploy the cloud routers. Each region is captured through a <code>cloudRegionName</code> child MO
<code>cloudtemplateExtNetwork</code>	Contains infra network configuration input that is external of the cloud. Contains a list of regions where cloud routers are configured for external networking. Each region is captured through a <code>cloudRegionName</code> child MO
<code>cloudtemplateVpnNetwork</code>	Contains information for setting up a VPN with an ACI on-premises site or another Cisco Cloud APIC site.


MO	Purpose
cloudtemplateIpSecTunnel	Captures the IP address of the IPsec peer in the ACI on-premises site.
cloudtemplateOspf	Captures the OSPF area to be used for the VPN connections.
cloudtemplateBgpEvpn	Captures the peer IP address, ASN, and so forth, for setting up the BGP session with the on-premises site.

In Cisco Cloud APIC, the layering of MOs is slightly different from a regular Cisco APIC due to the cloud template. In a regular Cisco APIC, you post logical MOs that go through two layers of translation:

1. Logical MO to resolved MO
2. Resolved MO to concrete MO

In Cisco Cloud APIC, there is an additional layer of translation for the infra network. This additional layer is where the cloud template translates logical MOs in the `cloudtemplate` namespace to logical MOs in the cloud namespace. For configurations outside of the infra network, you post logical MOs in the cloud namespace. In this case, the MOs go through the usual two-layer translation as in the regular Cisco APIC.

Figure 19: Cloud and Cloud Template MO Conversion



Note For information about configuring the cloud template, see [Configuring Cisco Cloud APIC Components, on page 45](#)


Managed Object Relations and Policy Resolution

Relationship-managed objects express the relation between managed object instances that do not share containment (parent-child) relations. MO relations are established between the source MO and a target MO in one of the following two ways:

- An explicit relation, such as with `cloudRsZoneAttach` and `cloudRsCloudEPgCtx`, defines a relationship that is based on the target MO distinguished name (DN).
- A named relation defines a relationship that is based on the target MO name.

The dotted lines in the following figure show several common MO relations.

Figure 20: MO Relations



For example, the dotted line between the cloud EPG and the VRF defines the relation between those two MOs. In this figure, the cloud EPG (`cloudEPg`) contains a relationship MO (`cloudRsCloudEPgCtx`) that is named with the name of the target VRF MO (`fvCtx`). For example, if production is the VRF name (`fvCtx.name=production`), then the relation name is `production` (`cloudRsCloudEPgCtx.tnFvCtxName=production`).

In the case of policy resolution based on named relations, if a target MO with a matching name is not found in the current tenant, the ACI cloud infrastructure tries to resolve in the common tenant. For example, if the user tenant cloud EPG contained a relationship MO targeted to a VRF that did not exist in the tenant, the system tries to resolve the relationship in the common tenant. If a named relation cannot be resolved in either the current tenant or the common tenant, the ACI cloud infrastructure attempts to resolve to a default policy. If a default policy exists in the current tenant, it is used. If it does not exist, the ACI cloud infrastructure looks for a default policy in the common tenant. Cloud context profile, VRF, and contract (security policy) named relations do not resolve to a default.

Default Policies



Warning Default policies can be modified or deleted. Deleting a default policy can result in a policy resolution process to complete abnormally.

The ACI cloud infrastructure includes default policies for many of its core functions. Examples of default policies include the following:

- Cloud AWS provider (for the infra tenant)
- Monitoring and statistics



Note To avoid confusion when implementing configurations that use default policies, document changes made to default policies. Be sure that there are no current or future configurations that rely on a default policy before deleting a default policy. For example, deleting a default firmware update policy could result in a problematic future firmware update.

A default policy serves multiple purposes:

- Allows a cloud infrastructure administrator to override the default values in the model.
- If an administrator does not provide an explicit policy, the Cisco CloudAPIC applies the default policy. An administrator can create a default policy and the Cisco Cloud APIC uses that unless the administrator provides any explicit policy.

The following scenarios describe common policy resolution behavior:

- A configuration explicitly refers to the default policy: if a default policy exists in the current tenant, it is used. Otherwise, the default policy in tenant **common** is used.
- A configuration refers to a named policy (not default) that does not exist in the current tenant or in tenant **common**: if the current tenant has a default policy, it is used. Otherwise, the default policy in tenant **common** is used.



Note The scenario above does not apply to a VRF in a tenant.

- A configuration does not refer to any policy name: if a default policy exists in the current tenant, it is used. Otherwise, the default policy in tenant **common** is used.

The policy model specifies that an object is using another policy by having a relation-managed object (MO) under that object and that relation MO refers to the target policy by name. If this relation does not explicitly refer to a policy by name, then the system tries to resolve a policy that is called default. Cloud context profiles and VRFs are exceptions to this rule.

Shared Services

Cloud EPGs in one tenant can communicate with cloud EPGs in another tenant through a contract interface that is contained in a shared tenant. Within the same tenant, a cloud EPG in one VRF can communicate with another cloud EPG in another VRF through a contract defined in the tenant. The contract interface is an MO that can be used as a contract consumption interface by the cloud EPGs that are contained in different tenants. By associating to an interface, a cloud EPG consumes the subjects that are represented by the interface to a contract contained in the shared tenant. Tenants can participate in a single contract, which is defined at some third place. More strict security requirements can be satisfied by defining the tenants, contract, subjects, and filter directions so that tenants remain isolated from one another.

Follow these guidelines when configuring shared services contracts:

- A shared service is supported only with non-overlapping and non-duplicate CIDR subnets. When configuring CIDR subnets for shared services, follow these guidelines:
 - CIDR subnets leaked from one VRF to another must be disjointed and must not overlap.
 - CIDR subnets advertised from multiple consumer networks into a VRF or vice versa must be disjointed and must not overlap.



CHAPTER 4

Configuring Cisco Cloud APIC Components

- [About Configuring the Cisco Cloud APIC, on page 45](#)
- [Configuring the Cisco Cloud APIC Using the GUI, on page 45](#)
- [Configuring Cisco Cloud APIC Using the REST API, on page 112](#)

About Configuring the Cisco Cloud APIC

You create the Cisco Cloud APIC components using either the Cisco Cloud APIC GUI or the REST API. This section explains how to create configuration, application management, operations, and administrative components.



Note

- For information about configuring a load balancer and service graph, see [Deploying Layer 4 to Layer 7 Services, on page 135](#).
- For information about the GUI, such as navigation and a list of configurable components, see [About the Cisco Cloud APIC GUI, on page 14](#).

Configuring the Cisco Cloud APIC Using the GUI

Creating a Tenant Using the Cisco Cloud APIC GUI For Release 4.2(2) and Earlier

This section explains how to create a tenant using the Cisco Cloud APIC GUI.

- Step 1** Click the **Intent** icon. The **Intent** menu appears.
- Step 2** Click the drop-down arrow below the **Intent** search box and choose **Application Management**.
A list of **Application Management** options appear in the **Intent** menu.
- Step 3** From the **Application Management** list in the **Intent** menu, click **Create Tenant**. The **Create Tenant** dialog box appears.
- Step 4** Enter the appropriate values in each field as listed in the following *Create Tenant Dialog Box Fields* table then continue.

Table 8: Create Tenant Dialog Box Fields

Properties	Description
Name	Enter the name of the tenant.
Description	Enter a description of the tenant.
Settings	
Add Security Domain	To add a security domain: <ol style="list-style-type: none"> Click Add Security Domain. The Select Security Domains dialog appears with a list of security domains in the left pane. Click to choose a security domain. Click Select to add the security domain to the tenant.
Trusted Tenant	Click to check (default) or uncheck the Enabled check box. Trusted Tenant is enabled when checked.
Cloud Account ID	Enter the cloud account ID.

Step 5 Click **Save** when finished.

Creating a Tenant Using the Cisco Cloud APIC GUI For Release 4.2(3) and Later

This section explains how to create a tenant using the Cisco Cloud APIC GUI.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Application Management**.

A list of **Application Management** options appear in the **Intent** menu.

Step 3 From the **Application Management** list in the **Intent** menu, click **Create Tenant**. The **Create Tenant** dialog box appears.

Step 4 Enter the appropriate values in each field as listed in the following *Create Tenant Dialog Box Fields* table then continue.

Table 9: Create Tenant Dialog Box Fields

Properties	Description
Name	Enter the name of the tenant.
Description	Enter a description of the tenant.
Settings	

Properties	Description
Add Security Domain	To add a security domain: <ol style="list-style-type: none"> Click Add Security Domain. The Select Security Domains dialog appears with a list of security domains in the left pane. Click to choose a security domain. Click Select to add the security domain to the tenant.
AWS Account ID	Enter the cloud account ID.
Access Type	Click to enable the tenant type: <ul style="list-style-type: none"> • Untrusted • Trusted • Organization

Step 5 Click **Save** when finished.

Configure a Tenant AWS Provider For Release 4.2(2) and Earlier

Before you begin

- AWS Provider is auto-configured for Infra tenant. You do not need to do anything to configure the AWS provider for the infra tenant.
- For all non-infra tenants, the AWS provider is configured either as a trusted tenant or as untrusted tenant. Our recommendation is to use trusted tenants because managing credentials is not easy. Also, each tenant must be in a separate AWS account. Sharing the same AWS account for multiple tenants is not allowed.

For a trusted tenant, establish the trust relationship first with the account in which Cisco Cloud APIC is deployed (the account for the infra tenant). To establish the trust relation and give all the required permissions to the Cisco Cloud APIC for accessing the tenant account, run the tenant role cloud-formation template in the tenant account. This template is available as a tenant-cft.json object in the S3 bucket that is named capic-common-[capicAccountId]-data in the infra tenant's AWS account. For security reasons, public access to this S3 bucket is not allowed, so the S3 bucket owner needs to download this file and use it in the tenant account.

- Untrusted tenants - use the account access and secret keys. The access and secret keys being used must be for an IAM user having these permissions at a minimum. The IAM role created must be named `ApicTenantRole`.



Note Cloud APIC does not disturb AWS resources created by other applications or users. It only manages the AWS resources created by itself.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "s3:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "events:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "logs:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "cloudtrail:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "cloudwatch:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "resource-groups:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "sqs:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": "elasticloadbalancing:*",
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "config:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": "iam:PassRole",
      "Resource": "arn:aws:iam::<account-id>:role/ApicTenantRole",

```

```

        "Effect": "Allow"
    }
}

```

- Add trust relationship:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "vpc-flow-logs.amazonaws.com",
        "AWS": "arn:aws:iam::<account-d>:root"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}

```

- Cloud APIC enforces ownership checks to prevent deployment of policies in the same tenant-region combination done either intentionally or by mistake. For example, assume that Cloud APIC is deployed in AWS account IA1 in region R1. Now you want to deploy a tenant TA1 in region R2. This tenant deployment i.e. account-region combination TA1-R2 is now owned by IA1-R1. If another Cloud APIC attempts to manage the same tenant-region combination later (say Capic2 in AWS account IA2 deployed in region R3), this will not be allowed because the current owner for the deployment TA1-R2 is IA1-R1. In other words, only one account in one region can be managed by one Cloud APIC. Example below shows some valid and wrong deployment combinations.

```

Capic1:
IA1-R1: TA1-R1 - ok
        TA1-R2 - ok

Capic2:
IA1-R2: TA1-R1 - not allowed
        TA1-R3 - ok

Capic3:
IA2-R1: TA1-R1 - not allowed
        TA1-R4 - ok
        TA2-R4 - ok

```

- Ownership enforcement is done using AWS Resource Groups. When a new tenant in account TA1 in region R2 is managed by Cloud APIC, a Resource Group CAPIC_TA1_R2 (e.g. CAPIC_123456789012_us-east-2) is created in the tenant account. This Resource Group has a resource tag AciOwnerTag with value IA1_R1_TA1_R2, assuming it was managed by Cloud APIC in account IA1 and deployed in region R1. If the AciOwnerTag mismatch happens, tenant-region management is aborted.

Here is a summary of AciOwnerTag mismatch cases:

- Initially Cloud APIC is installed in an account, and then taken down and Cloud APIC is installed in a different account. All existing tenant-region deployment will fail.
- Another Cloud APIC is managing the same tenant-region.

In ownership mismatch cases, **retry** (to setup tenant-region again) is not currently supported. As a workaround, if you are certain that no other Cloud APIC is managing the same tenant-region combination, logon to the tenant's AWS account and manually remove the affected Resource Group (e.g. CAPIC_123456789012_us-east-2). Next, reload Cloud APIC or delete and add the tenant again.

Step 1 In the Cloud APIC, configure the AWS Provider.

- a) On the **Intent** menu, choose **Tenants** > *tenant_name* from the drop-down.
- b) In the **Intent** pane, choose **Application Management** > *tenant_name* .

Step 2 Perform the following actions:

- a) Confirm there is a check in the **Trusted** Tenant checkbox.

The AWS account must be a Trusted account for the user tenant using the cloud.

- b) In the **Cloud Account ID** field, provide the Cloud account ID.
- c) Run the tenant role cloud-formation template available at the URL <https://capic-common-<infraAccountId>-data.s3.amazonaws.com/tenant-cft.json> which is in a s3 bucket in the infra tenant's AWS account.

Note Alternatively, keep the trusted flag unchecked and provide the access and secret keys as done normally for any tenant.

Step 3 Click **Save**.

Configuring a Tenant AWS Provider For Release 4.2(3) and Later

Before you begin

- AWS Provider is auto-configured for Infra tenant. You do not need to do anything to configure the AWS provider for the infra tenant.
- For all non-infra tenants, the AWS provider is configured either as a trusted tenant, untrusted tenant, or organization tenant. Our recommendation is to use trusted tenants because managing credentials is not easy. Also, each tenant must be in a separate AWS account. Sharing the same AWS account for multiple tenants is not allowed.

For a trusted tenant, establish the trust relationship first with the account in which Cisco Cloud APIC is deployed (the account for the infra tenant). To establish the trust relation and give all the required permissions to the Cisco Cloud APIC for accessing the tenant account, first create a tenant and assign the Trusted tag to that tenant as the Access Type. Then, bring up that new trusted tenant again by clicking on the tenant name in the Tenants page, and in the AWS Account area in the tenant window, click the Run the CloudFormation template link.

- Organization tenants are for adding tenant accounts that are part of the organization. This requires deploying the Cisco Cloud APIC in the master account of the organization.
- Untrusted tenants use the account access and secret keys. The access and secret keys being used must be for an IAM user having these permissions at a minimum. The IAM role created must be named `ApicTenantRole`.



Note Cloud APIC does not disturb AWS resources created by other applications or users. It only manages the AWS resources created by itself.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "s3:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "events:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "logs:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "cloudtrail:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "cloudwatch:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "resource-groups:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "sqs:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": "elasticloadbalancing:*",
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "config:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": "iam:PassRole",
      "Resource": "arn:aws:iam::<account-id>:role/ApicTenantRole",

```

```

        "Effect": "Allow"
      }
    ]
  }
}

```

- Add trust relationship:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "vpc-flow-logs.amazonaws.com",
        "AWS": "arn:aws:iam:<infra-account-id>:root"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}

```

- The Cloud APIC uses the OrganizationAccountAccessRole IAM role to manage policies for AWS Organization tenants.
 - If you created an AWS account within the existing organization in the master account, the OrganizationAccountAccessRole IAM role is automatically assigned to that created AWS account. You do not have to manually configure the OrganizationAccountAccessRole IAM role in AWS in this case.
 - If the master account invited an existing AWS account to join the organization, then you must manually configure the OrganizationAccountAccessRole IAM role in AWS. Configure the OrganizationAccountAccessRole IAM role in AWS for the organization tenant and verify that it has Cloud APIC-related permissions available.

The OrganizationAccountAccessRole IAM role, together with the SCP (Service Control Policy) used for the organization or the account, must have the minimum permissions that are required by the Cloud APIC to manage policies for the tenants. The access policy requirement is the same as the requirement for the trusted or untrusted tenants.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "s3:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "events:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }, {
      "Action": [
        "logs:*"
      ]
    }
  ]
}

```

```

    ],
    "Resource": "*",
    "Effect": "Allow"
  }, {
    "Action": [
      "cloudtrail:*"
    ],
    "Resource": "*",
    "Effect": "Allow"
  }, {
    "Action": [
      "cloudwatch:*"
    ],
    "Resource": "*",
    "Effect": "Allow"
  }, {
    "Action": [
      "resource-groups:*"
    ],
    "Resource": "*",
    "Effect": "Allow"
  }, {
    "Action": [
      "sqs:*"
    ],
    "Resource": "*",
    "Effect": "Allow"
  }, {
    "Action": "elasticloadbalancing:*",
    "Resource": "*",
    "Effect": "Allow"
  }, {
    "Action": [
      "config:*"
    ],
    "Resource": "*",
    "Effect": "Allow"
  }, {
    "Action": "iam:PassRole",
    "Resource": "*",
    "Effect": "Allow"
  }
}
]
}

```

To add a trust relationship for an Organization tenant:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "vpc-flow-logs.amazonaws.com",
        "AWS": "arn:aws:iam::<infra-account-id>:root"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}

```

- Cloud APIC enforces ownership checks to prevent deployment of policies in the same tenant-region combination done either intentionally or by mistake. For example, assume that Cloud APIC is deployed

in AWS account IA1 in region R1. Now you want to deploy a tenant TA1 in region R2. This tenant deployment i.e. account-region combination TA1-R2 is now owned by IA1-R1. If another Cloud APIC attempts to manage the same tenant-region combination later (say Capic2 in AWS account IA2 deployed in region R3), this will not be allowed because the current owner for the deployment TA1-R2 is IA1-R1. In other words, only one account in one region can be managed by one Cloud APIC. Example below shows some valid and wrong deployment combinations.

```
Capic1:
IA1-R1: TA1-R1 - ok
        TA1-R2 - ok

Capic2:
IA1-R2: TA1-R1 - not allowed
        TA1-R3 - ok

Capic3:
IA2-R1: TA1-R1 - not allowed
        TA1-R4 - ok
        TA2-R4 - ok
```

- Ownership enforcement is done using AWS Resource Groups. When a new tenant in account TA1 in region R2 is managed by Cloud APIC, a Resource Group CAPIC_TA1_R2 (e.g. CAPIC_123456789012_us-east-2) is created in the tenant account. This Resource Group has a resource tag AciOwnerTag with value IA1_R1_TA1_R2, assuming it was managed by Cloud APIC in account IA1 and deployed in region R1. If the AciOwnerTag mismatch happens, tenant-region management is aborted.

Here is a summary of AciOwnerTag mismatch cases:

- Initially Cloud APIC is installed in an account, and then taken down and Cloud APIC is installed in a different account. All existing tenant-region deployment will fail.
- Another Cloud APIC is managing the same tenant-region.

In ownership mismatch cases, **retry** (to setup tenant-region again) is not currently supported. As a workaround, if you are certain that no other Cloud APIC is managing the same tenant-region combination, logon to the tenant's AWS account and manually remove the affected Resource Group (e.g. CAPIC_123456789012_us-east-2). Next, reload Cloud APIC or delete and add the tenant again.

Step 1 In the Cloud APIC, configure the AWS Provider.

- On the **Intent** menu, choose **Tenants** > *tenant_name* from the drop-down.
- In the **Intent** pane, choose **Application Management** > *tenant_name* .

Step 2 Perform the following actions:

- In the **AWS Account ID** field, provide the cloud account ID.
- In the **Access Type** area, choose **Trusted**.

The AWS account must be a Trusted account for the user tenant that is using the cloud.

- Click **Save**.
- Bring up the new trusted tenant again by clicking on the tenant name in the **Tenants** page.

In the **AWS Account** area in the tenant **Overview** page, you will see the following message: "In order to deploy any configuration from this tenant, you must create a trusted role in the tenant AWS account which will establish trust with the AWS infra account. To do so, open the link below to run the CloudFormation template."

- e) Click the **Run the CloudFormation** template link.

This returns you to the AWS sign in page, which should be pre-populated with the necessary AWS account information that you entered earlier in these procedures in the Cloud APIC GUI.

- f) Click **Next** in the AWS sign in page after verifying that the sign-in information is correct.
g) Run the tenant role cloud-formation template in the tenant account.

Note Alternatively, keep the trusted flag unchecked and provide the access and secret keys as done normally for any tenant.

Step 3 Click **Save**.

Creating an Application Profile Using the Cisco Cloud APIC GUI

This section explains how to create an application profile using the Cisco Cloud APIC GUI.

Before you begin

Create a tenant.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Application Management**.

A list of **Application Management** options appear in the **Intent** menu.

Step 3 From the **Application Management** list in the **Intent** menu, click **Create Application Profile**. The **Create Application Profile** dialog box appears.

Step 4 Enter a name in the **Name** field.

Step 5 Choose a tenant:

- a) Click **Select Tenant**.

The **Select Tenant** dialog box appears.

- b) From the **Select Tenant** dialog, click to choose a tenant in the left column then click **Select**.

You return to the **Create Application Profile** dialog box.

Step 6 Enter a description in the **Description** field.

Step 7 Click **Save** when finished.

Creating a VRF Using the Cisco Cloud APIC GUI

This section explains how to create a VRF using the Cisco Cloud APIC GUI.

Before you begin

Create a tenant.

- Step 1** Click the **Intent** icon. The **Intent** menu appears.
- Step 2** Click the drop-down arrow below the **Intent** search box and choose **Application Management**.
A list of **Application Management** options appear in the **Intent** menu.
- Step 3** From the **Application Management** list in the **Intent** menu, click **Create VRF**. The **Create VRF** dialog box appears.
- Step 4** Enter the appropriate values in each field as listed in the following *Create VRF Dialog Box Fields* table then continue.

Table 10: Create VRF Dialog Box Fields

Properties	Description
General	
Name	Enter a name for the VRF in the Name field. All VRFs are assigned a <i>vrfEncoded</i> value. If the Tenant and VRF name combination has more than 32 characters, then, a VRF name (which also contains the tenant name) is identified in the cloud router using the <i>vrfEncoded</i> value. To see the <i>vrfEncoded</i> value, navigate to Application Management > VRFs subtab. Click a VRF on the right hand pane and look for <i>Encoded VRF Name in Cloud Router</i> .
Tenant	To choose a tenant: a. Click Select Tenant . The Select Tenant dialog box appears. b. From the Select Tenant dialog, click to choose a tenant in the left column then click Select . You return to the Create VRF dialog box.
Description	Enter a description of the VRF.
Settings > IPv4 unicast address family BGP targets	
Add Filter	a. Click the Add Route Target option for the unicast address family BGP target you want to configure. b. Click to choose the following options for the Type field: <ul style="list-style-type: none"> • Export—The route target can be exported to other VRFs • Import—The route target is imported from other VRFs • Enter the route target that can be exported from the current VRF or imported into the current VRF in the Route Target text box.

Step 5 When finished, click **Save**.

Creating an External Network Using the Cisco Cloud APIC GUI

This procedure describes how to create an external network. You can have a single external network that can connect to multiple routers on the on-premises site, or you can have multiple external networks with multiple VRFs that you can use to connect to CCRs.

Before you begin

You must have a hub network created before you can create an external network.

- Step 1** In the left navigation bar, navigate to **Application Management > External Networks**.
The configured external networks are displayed.
- Step 2** Click **Actions**, then choose **Create External Network**.
The **Create External Network** window appears.
- Step 3** Enter the appropriate values in each field as listed in the following *Create External Network Dialog Box Fields* table then continue.

Table 11: Create External Network Dialog Box Fields

Properties	Description
General	
Name	Enter the name for the external network.
VRF	<p>This external VRF will be used for external connectivity with external non-ACI devices. You can create multiple external VRFs for this purpose.</p> <p>This VRF will be identified as an external VRF if the VRF has all three of the following characteristics:</p> <ul style="list-style-type: none"> • Configured under the infra tenant • Associated with an external network • Not associated with a cloud context profile <p>Any VRF that is associated with an external network becomes an external VRF. The external VRF is not allowed to be associated with a cloud context profile or subnet.</p> <p>To choose an external VRF:</p> <ol style="list-style-type: none"> a. Click Select VRF. The Select VRF dialog box appears. b. From the Select VRF dialog, click to choose a VRF in the left column. You can also create a VRF using the + Create VRF option. c. Click Select. You return to the Create External Network dialog box.

Properties	Description
Router Type	Choose the router type: <ul style="list-style-type: none"> • CCR: <ul style="list-style-type: none"> • For releases prior to 25.0(3), the Cisco Cloud Services Router 1000V • For release 25.0(3) and later, the Cisco Catalyst 8000V • TGW: An AWS transit gateway router
Host Router Name	This field appears if you select CCR as the Router Type . This field is not editable. The default host router is automatically selected.
Hub Network	This field appears if you select TGW as the Router Type . To choose a hub network: <ol style="list-style-type: none"> a. Click Select Hub Network. The Select Hub Network dialog box appears. b. In the Select Hub Network dialog box, click the desired hub network from the list and then click Select. You are returned to the Create External Network page.
Settings	
Regions	To choose a region: <ol style="list-style-type: none"> a. Click Add Regions. The Select Regions dialog box appears. The regions that you selected as part of the First Time Setup are displayed here. b. From the Select Regions dialog, click to choose a region in the left column then click Select. You return to the Create External Network dialog box.

Properties	Description
VPN Networks	

Properties	Description
	<p>The VPN networks entries are used for external connectivity. All configured VPN networks will be applied to all the selected regions.</p> <p>To add a VPN network:</p> <ol style="list-style-type: none"> Click Add VPN Network. The Add VPN Network dialog box appears. In the Name field, enter a name for the VPN network. Click + Add IPsec Peer. The Add IPsec Tunnel Destination window appears. Enter values for the following fields for the IPsec tunnel destination that you want to add: <ul style="list-style-type: none"> Public IP of IPsec Tunnel Peer Pre-Shared Key IKE Version: Select ikev1 or ikev2 for IPsec tunnel connectivity BGP Peer ASN Subnet Pool Name: Click Select Subnet Pool Name. The Select Subnet Pool Name dialog box appears. Select one of the available subnet pools that are listed, then click Select. <p>Note Additional IPsec tunnel subnet pools can be added in the External Networks page, or through the Cloud APIC First Time Set Up, if necessary. For more information on adding additional subnet pools through the Cloud APIC First Time Set Up, see the chapter "Configuring Cisco Cloud APIC Using the Setup Wizard" in the Cisco Cloud APIC for AWS Installation Guide, Release 25.0(1)-25.0(4) and later. The subnet pool size should be large enough to accommodate the number of IPsec tunnels that will get created.</p> <ul style="list-style-type: none"> IPsec Tunnel Source Interfaces: Using the entries in this field, the Cisco Cloud APIC creates one IPsec tunnel from each selected source interface to the destination IP address. <p>Note ikev2 is the default option in this field. The IPsec tunnel source interfaces feature is supported only with the IKEv2 configuration.</p> <p>gig3 is selected by default. Choose one or more from the following interfaces:</p> <ul style="list-style-type: none"> gig2: The GigabitEthernet2 interface gig3: The GigabitEthernet3 interface gig4: The GigabitEthernet4 interface <p>Note After you have configured the IPsec tunnel source interfaces in this external network, you can then configure IPsec tunnel source interfaces in additional networks where tunnels to the same destination can be formed, as described in Routing Policies: Release 25.0(2), on page 9.</p>

Properties	Description
	<p>e. Click Add to add this IPsec tunnel destination.</p> <p>You return to the Add VPN Network window.</p> <p>Click + Add IPsec Peer if you want to add another IPsec tunnel destination.</p> <p>f. Click Add in the Add VPN Network dialog box.</p> <p>You return to the Create External Network dialog box.</p>

Step 4 When you have finished creating the external network, click **Save**.
After you click **Save** in the **Create External Network** window, cloud routers are then configured in AWS.

Configuring the Global Inter-VRF Route Leak Policy

The global inter-VRF route leak policy feature is introduced in release 25.0(2).

Before you begin

Review the information provided in [Global Inter-VRF Route Leak Policy, on page 10](#) before making any changes in the **Contract Based Routing** area in the **Cloud APIC Setup** window.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Configuration**.

A list of options appear in the **Intent** menu.

Step 3 From the **Configuration** list in the **Intent** menu, click **Cloud APIC Setup**.
The **Set up - Overview** dialog box appears.

Step 4 In the **Contract Based Routing** area, note the current setting for the **Contract Based Routing** field.

The **Contract Based Routing** setting reflects the current internal VRF route leak policy, which is a global policy under the infra tenant where a Boolean flag is used to indicate whether contracts can drive routes in the absence of route maps:

- **Off**: Default setting. Indicates that routes are not leaked based on contracts, and are leaked based on route maps instead.
- **On**: Indicates that routes are leaked based on contracts in the absence of route maps. When enabled, contracts drive routing when route maps are not configured. When route maps exist, route maps always drives routing.

Step 5 Determine if you want to change the current setting for the **Contract Based Routing** field.

Follow these procedures if you toggle from one setting to another:

- **Toggle from On setting to Off (disabling contract-based routing)**: In this situation, the assumption is that you have contract-based routing configured currently and you want to toggle over to route map-based routing. This can be disruptive if the route map-based routing is not configured before you toggle from contract-based routing to route map-based routing.

Before toggling from the **On** setting to the **Off** setting in this situation, make the following changes:

- a. Between all pairs of VRFs that have existing contracts, enable route map-based route leaking.
Follow the procedures provided in [Configuring Leak Routes Using the Cisco Cloud APIC GUI, on page 62](#).
- b. Disable the contract-based route policy in the global policy.
Toggle the switch in the **Contract Based Routing** field from the **On** setting to the **Off** setting to toggle from contract-based routing to route map-based routing.
- c. Change the routing to reflect any granularity that is required based on the new route map-based routing that you enabled.

- **Toggling from Off setting to On (enabling contract-based routing):** In this situation, the assumption is that you have route map-based routing configured currently and you want to toggle over to contract-based routing. This is not a disruptive operation, but rather is an additive operation, since both contracts and route maps can be enabled between a pair of VRFs. In that situation, route maps take precedence over contracts when enabling routing. With route map-based routing enabled, adding contract-based routing should be non-disruptive.

For that reason, you do not have to make any changes before toggling from the **Off** setting to the **On** setting in this situation. However, if you do not want to have both contracts and route maps enabled between a pair of VRFs, and you want to move completely to contract-based routing, you should completely set up contracts between the VRFs and delete the route maps between the VRFs before toggling to the **On** setting in the **Contract Based Routing** field.

Step 6 If you want to change the current setting for the **Contract Based Routing** area, toggle the setting based on the type of routing that you want.

Step 7 Click **Done** when you have finished the **Cloud APIC Setup** configurations.

Configuring Leak Routes Using the Cisco Cloud APIC GUI

The procedures for configuring leak routes using the Cisco Cloud APIC GUI will vary slightly, depending on the release:

- For releases prior to 25.0(2), you can configure an independent routing policy to specify which routes to leak between internal and external VRFs when you are setting up routing between an ACI cloud site and an external destination using the external connectivity feature. See [Configuring Inter-VRF Route Leaking Using the Cisco Cloud APIC GUI, on page 62](#) for those procedures.
- For releases 25.0(2) and later, support is available for route maps-based route leaking between a pair of internal VRFs. See [Configuring Leak Routes for Internal VRFs Using the Cisco Cloud APIC GUI, on page 65](#) for those procedures.

Configuring Inter-VRF Route Leaking Using the Cisco Cloud APIC GUI

Configuring leak routes is part of the release 25.0(1) update where routing and security policies are configured separately. Using inter-VRF routing, you can configure an independent routing policy to specify which routes to leak between internal and external VRFs when you are setting up routing between an ACI cloud site and an external destination using the external connectivity feature. See [Understanding Supported Routing and Security Policies, on page 7](#) for more information.

The external destination must be configured manually using the [Enabling Connectivity From the AWS Site to External Devices, on page 67](#) procedures. The external destination could be another cloud site, an ACI on-premises site or a branch office.



Note

- Use these procedures to configure routing policies independent of security policies only between internal and external VRFs, based on updates provided in release 25.0(1).
- Do not use these procedures to configure routing between a pair of internal VRFs; use contracts as you normally would prior to release 25.0(1) in that case.

- Step 1** In the left navigation bar, navigate to **Application Management > VRFs**.
The configured VRFs are displayed.
- Step 2** Click the **Leak Routes** tab.
Any already-configured leak routes are displayed.
- Step 3** Click **Actions**, then choose **Create Leak Route**.
The **Create Leak Route** window appears.
- Step 4** Enter the appropriate values in each field as listed in the following *Create Leak Routes Dialog Box Fields* table then continue.

Table 12: Create Leak Routes Dialog Box Fields

Properties	Description
Source VRF	To choose a source VRF: <ol style="list-style-type: none"> a. Click Select a Source VRF. The Select a VRF dialog box appears. b. From the Select a VRF dialog, click to choose a VRF in the left column to use for the source VRF. Note that the source VRF can be an internal or an external VRF. c. Click Select to select this source VRF. You return to the Create Leak Route dialog box.
Destination VRF	To choose a destination VRF: <ol style="list-style-type: none"> a. Click Select a Destination VRF. The Select a VRF dialog box appears. b. From the Select a VRF dialog, click to choose a VRF in the left column to use for the destination VRF. Note that the destination VRF cannot be an internal VRF if the source VRF is also internal VRF. c. Click Select to select this destination VRF. You return to the Create Leak Route dialog box.

Properties	Description
Type	<p>Choose the type of leaked route that you want to configure:</p> <ul style="list-style-type: none"> • Leak All: Select to configure all routes to leak from the source VRF to the destination VRF. The entry 0.0.0.0/0 is entered automatically in the subnet IP area by default in this case. • Subnet IP: Select to configure a specific subnet IP address as the route to leak from the source VRF to the destination VRF. The Subnet IP box appears. In the Subnet IP box, enter a subnet IP address as the route to leak between VRFs.

Step 5 When finished, click **Save**.
The **Success** window appears.

Step 6 Determine if you want to configure additional inter-VRF route leaking.

- If you want to add another route to leak between a pair of VRFs, click the **Add Another Leak Route** option in the **Success** window.

You are returned to the **Add Leak Route** window. Repeat [Step 4, on page 63](#) through [Step 5, on page 64](#) to configure another route to leak between a pair of VRFs.

- If you want to add a reverse route, where:

- The destination VRF from the previous configuration now becomes the source VRF, and
- The source VRF from the previous configuration now becomes the destination VRF

Then click the **Add Reverse Leak Route** option in the **Success** window.

You are returned to the **Add Leak Route** window. Repeat [Step 4, on page 63](#) through [Step 5, on page 64](#) to configure another route, but this time:

- In the **Source VRF** field, select the VRF that you had selected as a destination VRF in the previous configuration.
- In the **Destination VRF** field, select the VRF that you had selected as a source VRF in the previous configuration.

Step 7 When you have finished configuring leak routes, click **Done**.

The **Leak Routes** tab in the main **VRFs** page is displayed again, with the newly configured leak route displayed.

Step 8 To get more information on a source or destination VRF, or to make changes to a configured leak route, double-click the VRF in the **Leak Routes** tab in the main **VRFs** page.
The **Overview** page for that VRF is displayed.

Step 9 Click the **Application Management** tab at the top of the VRF page, then click the **Leak Routes** tab in the left nav bar. The leak routes associated with this particular VRF are displayed.

Step 10 Configure additional leak routes associated with this VRF, if necessary.

- To add a leak route from this VRF, click **Actions**, then choose **Add Leak Route from <VRF_name>**.

The **Add Leak Route** window appears. Enter the necessary information as you did previously using the information in [Step 4, on page 63](#). Note that the entry in the **Source VRF** is pre-selected and cannot be changed in this situation.

- To add a leak route to this VRF, click **Actions**, then choose **Add Leak Route to <VRF_name>**.

The **Add Leak Route** window appears. Enter the necessary information as you did previously using the information in [Step 4, on page 63](#). Note that the entry in the **Destination VRF** is pre-selected and cannot be changed in this situation.

What to do next

You have now configured the routing policy. Since the routing and security policies are separate, you now need to configure the security policy separately:

- [Creating an EPG Using the Cisco Cloud APIC GUI, on page 71](#): Use these procedures to create an external EPG.
- [Creating a Contract Using the Cisco Cloud APIC GUI, on page 76](#): Use these procedures to create a contract between the external EPG and the cloud EPG.

Configuring Leak Routes for Internal VRFs Using the Cisco Cloud APIC GUI

Beginning with release 25.0(2), support is available for route maps-based route leaking between a pair of internal VRFs, as described in [Route Leaking Between Internal VRFs, on page 9](#). This feature is an extension of the routing and security split update provided in release 25.0(1), where routing and security policies are configured separately.

- Step 1** In the left navigation bar, navigate to **Application Management > VRFs**.
The configured VRFs are displayed.
- Step 2** Click the **Leak Routes** tab.
Any already-configured leak routes are displayed.
- Step 3** Click **Actions**, then choose **Create Leak Route**.
The **Create Leak Route** window appears.
- Step 4** Enter the appropriate values in each field as listed in the following *Create Leak Routes Dialog Box Fields* table then continue.

Table 13: Create Leak Routes Dialog Box Fields

Properties	Description
Source VRF	<p>To choose a source VRF:</p> <ol style="list-style-type: none"> Click Select a Source VRF. The Select a VRF dialog box appears. From the Select a VRF dialog, click to choose a VRF in the left column to use for the source VRF. Because this procedure is for route maps-based route leaking between a pair of internal VRFs, choose an internal VRF for the source VRF. Click Select to select this source VRF. You return to the Create Leak Route dialog box.

Properties	Description
Destination VRF	<p>To choose a destination VRF:</p> <ol style="list-style-type: none"> a. Click Select a Destination VRF. The Select a VRF dialog box appears. b. From the Select a VRF dialog, click to choose a VRF in the left column to use for the destination VRF. Because this procedure is for route maps-based route leaking between a pair of internal VRFs, choose an internal VRF for the destination VRF. c. Click Select to select this destination VRF. You return to the Create Leak Route dialog box.
Type	<p>Choose the type of leaked route that you want to configure:</p> <ul style="list-style-type: none"> • Leak All: Select to configure all routes to leak from the source VRF to the destination VRF. The entry <code>0.0.0.0/0</code> is entered automatically in the subnet IP area by default in this case. • Subnet IP: Select to configure a specific subnet IP address as the route to leak from the source VRF to the destination VRF. The Subnet IP box appears. In the Subnet IP box, enter a subnet IP address as the route to leak between VRFs.

Step 5 When finished, click **Save**.

The **Success** window appears.

Step 6 Determine if you want to configure additional inter-VRF route leaking.

- If you want to add another route to leak between a pair of VRFs, click the **Add Another Leak Route** option in the **Success** window.

You are returned to the **Add Leak Route** window. Repeat [Step 4, on page 65](#) through [Step 5, on page 66](#) to configure another route to leak between a pair of VRFs.

- If you want to add a reverse route, where:
 - The destination VRF from the previous configuration now becomes the source VRF, and
 - The source VRF from the previous configuration now becomes the destination VRF

Then click the **Add Reverse Leak Route** option in the **Success** window.

You are returned to the **Add Leak Route** window. Repeat [Step 4, on page 65](#) through [Step 5, on page 66](#) to configure another route, but this time:

- In the **Source VRF** field, select the VRF that you had selected as a destination VRF in the previous configuration.
- In the **Destination VRF** field, select the VRF that you had selected as a source VRF in the previous configuration.

Step 7 When you have finished configuring leak routes, click **Done**.

The **Leak Routes** tab in the main **VRFs** page is displayed again, with the newly configured leak route displayed.

- Step 8** To get more information on a source or destination VRF, or to make changes to a configured leak route, double-click the VRF in the **Leak Routes** tab in the main **VRFs** page. The **Overview** page for that VRF is displayed.
- Step 9** Click the **Application Management** tab at the top of the VRF page, then click the **Leak Routes** tab in the left nav bar. The leak routes associated with this particular VRF are displayed.
- Step 10** Configure additional leak routes associated with this VRF, if necessary.
- To add a leak route from this VRF, click **Actions**, then choose **Add Leak Route from <VRF_name>**.
The **Add Leak Route** window appears. Enter the necessary information as you did previously using the information in [Step 4, on page 65](#). Note that the entry in the **Source VRF** is pre-selected and cannot be changed in this situation.
 - To add a leak route to this VRF, click **Actions**, then choose **Add Leak Route to <VRF_name>**.
The **Add Leak Route** window appears. Enter the necessary information as you did previously using the information in [Step 4, on page 65](#). Note that the entry in the **Destination VRF** is pre-selected and cannot be changed in this situation.

Enabling Connectivity From the AWS Site to External Devices

Follow these procedures to manually enable IPv4 connectivity from the infra VPC CCRs to any external device with IPsec/BGP.

Downloading the External Device Configuration Files

- Step 1** In the Cisco Cloud APIC GUI, click on **Dashboard**.
The **Dashboard** view for the Cisco Cloud APIC appears.
- Step 2** Navigate to **Infrastructure > External Connectivity**.
The **External Connectivity** window appears.
- Step 3** Click **Actions > Download External Device Configuration Files**.
The **Download External Device Configuration Files** pop-up appears.
- Step 4** Select the external device configuration files to download and click **Download**.
This action downloads a zip file that contains configuration information that you will use to manually configure the external device for IPv4 connectivity to the CCRs.

Enabling Connectivity From the AWS Site to External Devices

- Step 1** Gather the necessary information that you will need to manually enable IPv4 connectivity from the infra VPC CCRs to any external device without EVPN.
- Step 2** Log into the external device.
- Step 3** Enter the configuration information to connect an external networking device.

If you downloaded the external device configuration files using the instructions in [Downloading the External Device Configuration Files, on page 67](#), locate the configuration information for the first tunnel and enter that configuration information.

Following is an example of what the external device configuration file might look like for the first tunnel:

```
! The following file contains configuration recommendation to connect an external networking device
with the cloud ACI Fabric
! The configurations here are provided for an IOS-XE based device. The user is expected to understand
the configs and make any necessary amends before using them
! on the external device. Cisco does not assume any responsibility for the correctness of the config.

! Tunnel to 128.107.72.122 1.100 [ikev2] for
hctunnIf.acct-[infra]/region-[westus]/context-[overlay-1]-addr-[10.115.9.128/25]/csr-[ct_routerp_westus_0:0]/tunn-34
! USER-DEFINED: please define gig-gateway: GIG-GATEWAY
! USER-DEFINED: please define GigabitEthernet2 if required
! USER-DEFINED: please define tunnel-id: 100 if required
! USER-DEFINED: please define vrf-name: infra:externalvrf1 if required
! USER-DEFINED: please define gig3-public-ip: 13.88.168.176 if 0.0.0.0 ip still not provided by AWS.
! Device:          128.107.72.122
! Tunnel ID:       100
! Tunnel counter:  1
! Tunnel address:  5.16.1.9
! Tunnel Dn:
acct-[infra]/region-[westus]/context-[overlay-1]-addr-[10.115.9.128/25]/csr-[ct_routerp_westus_0:0]/tunn-34
! VRF name:        infra:externalvrf1
! ikev:            ikev2
! Bgp Peer addr:   5.16.1.10
! Bgp Peer asn:    65015
! Gig3 Public ip:  13.88.168.176
! PreShared key:   devicelazure
! ikev profile name: ikev2-100

vrf definition infra:externalvrf1
  rd 1:1

  address-family ipv4
    route-target export 64550:1
    route-target import 64550:1
  exit-address-family
exit

crypto ikev2 proposal ikev2-infra:externalvrf1
  encryption aes-cbc-256 aes-cbc-192 aes-cbc-128
  integrity sha512 sha384 sha256 sha1
  group 24 21 20 19 16 15 14 2
exit

crypto ikev2 policy ikev2-infra:externalvrf1
  proposal ikev2-infra:externalvrf1
exit

crypto ikev2 keyring keyring-ikev2-100
  peer peer-ikev2-keyring
    address 13.88.168.176
    pre-shared-key devicelazure
  exit
exit

crypto ikev2 profile ikev2-100
  match address local interface GigabitEthernet2
  match identity remote address 13.88.168.176 255.255.255.255
  identity local address 128.107.72.122
```

```
    authentication remote pre-share
    authentication local pre-share
    keyring local keyring-ikev2-100
    lifetime 3600
    dpd 10 5 on-demand
exit

crypto ipsec transform-set ikev2-100 esp-gcm 256
    mode tunnel
exit

crypto ipsec profile ikev2-100
    set transform-set ikev2-100
    set pfs group14
    set ikev2-profile ikev2-100
exit

interface Tunnel100
    vrf forwarding infra:externalvrf1
    ip address 5.16.1.10 255.255.255.252
    ip mtu 1400
    ip tcp adjust-mss 1400
    tunnel source GigabitEthernet2
    tunnel mode ipsec ipv4
    tunnel destination 13.88.168.176
    tunnel protection ipsec profile ikev2-100
exit

ip route 13.88.168.176 255.255.255.255 GigabitEthernet2 GIG-GATEWAY

router bgp 65015

address-family ipv4 vrf infra:externalvrf1
    redistribute connected
    maximum-paths eibgp 32

    neighbor 5.16.1.9 remote-as 65008
    neighbor 5.16.1.9 ebgp-multihop 255
    neighbor 5.16.1.9 activate
    neighbor 5.16.1.9 send-community both

    distance bgp 20 200 20
exit-address-family
```

The following figures provide more information on what each set of fields is used for in the external device configuration file:

- The fields shown in the following figure are used to configure these areas:
 - VRF definition
 - IPSec global configurations

```

vrf definition Ext-V1
rd 1:10
!
address-family ipv4
  route-target export 64550:10
  route-target import 64550:10
!
crypto isakmp policy 10
encryption aes
authentication pre-share
group 2
lifetime 28800
!
crypto isakmp keepalive 10 10 periodic
crypto isakmp aggressive-mode disable
!

```

VRF Definition

IPSec Global Configurations

- The fields shown in the following figure are used to configure these areas:

- IPSec and ikev1 per tunnel configurations
- BGP configurations for the VRF neighbor

```

!
crypto keyring Ext-V1-1000-ike
  pre-shared-key address <50.18.55.126>[cAPIC CSR Gig3 Public IP] key <abodef12345>
!
crypto isakmp profile Ext-V1-1000-ike
  keyring Ext-V1-1000-ike
  match identity address <50.18.55.126>[cAPIC CSR1 gig3 Public IP] 255.255.255.255
!
crypto ipsec transform-set Ext-V1-1000-ike esp-aes esp-sha-hmac
mode tunnel
!
crypto ipsec profile Ext-V1-1000-ike
set security-association lifetime kilobytes disable
set security-association replay window-size 512
set transform-set Ext-V1-1000-ike
set pfs group14
!
interface Tunnel1000
vrf forwarding Ext-V1
ip address 50.50.0.2[cAPIC CSR BGP Peer Addr] 255.255.255.252
ip mtu 1400
ip tcp adjust-mss 1400
tunnel source GigabitEthernet2
tunnel mode ipsec ipv4
tunnel destination <50.18.55.126>[cAPIC CSR1 gig3 Public IP]
tunnel protection ipsec profile Ext-V1-1000-ike
!
router bgp 64550
!
address-family ipv4 vrf Ext-V1
 redistribute connected
 neighbor <50.50.0.1>[cAPIC CSR1 Tunnel Inner IP Addr] remote-as 1234
 neighbor 50.50.0.1 ebgp-multihop 255
 neighbor 50.50.0.1 activate
 neighbor 50.50.0.1 send-community both
 neighbor <50.50.0.5>[cAPIC CSR1 Tunnel Inner IP Addr] remote-as 1234
 neighbor 50.50.0.5 ebgp-multihop 255
 neighbor 50.50.0.5 activate
 neighbor 50.50.0.5 send-community both
 distance bgp 20 200 20
!
ip route 50.18.55.126[cAPIC CSR1 gig3 Public IP] 255.255.255.255 GigabitEthernet2 10.10.0.103
!

```

IPSec and ikev1
Per Tunnel Configurations

BGP Configurations for VRF Neighbor

- The fields shown in the following figure are used to configure these areas:

- Ikev2 global configurations
- IPSec and ikev2 per tunnel configurations


```
crypto ikev2 proposal ikev2-1
  encryption aes-cbc-256 aes-cbc-192 aes-cbc-128
  integrity sha512 sha384 sha256 sha1
  group 24 21 20 19 16 15 14 2
  !
crypto ikev2 policy ikev2-1
  proposal ikev2-1
  !
crypto ikev2 keyring ikev2-2000
  peer peer-ikev2-keyring
  address 35.81.94.248 [cAPIC CSR1 gig3 Public IP]
  pre-shared-key abcdefg12345
  !
crypto ikev2 profile ikev2-2000
  match address local interface GigabitEthernet3
  match identity remote address 35.81.94.248[cAPIC CSR1 gig3 Public IP] 255.255.255.255
  identity local address 52.53.49.193 [Local Device tunnel source interface Public IP (Gig3 public IP)]
  authentication remote pre-share
  authentication local pre-share
  keyring local keyring-ikev2-2000
  lifetime 3600
  dpd 10 5 on-demand
  !
crypto ipsec transform-set ikev2-2000 esp-gcm 256
  mode tunnel
  !
crypto ipsec profile ikev2-2000
  set transform-set ikev2-2000
  set pfs group14
  set ikev2-profile ikev2-2000
  !
interface Tunnel2000
  vrf forwarding Ext-V1
  ip address 50.50.0.14 [cAPIC CSR1 BGP Peer Addr] 255.255.255.252
  ip mtu 1400
  ip tcp adjust-mss 1400
  tunnel source GigabitEthernet3
  tunnel mode ipsec ipv4
  tunnel destination 35.81.94.248[cAPIC CSR1 gig3 Public IP]
  tunnel protection ipsec profile ikev2-2000
```

Ikev2 Global Configurations

IPSec and Ikev2
Per Tunnel Configurations

Step 4 Repeat the previous step to configure additional tunnels.

Creating an EPG Using the Cisco Cloud APIC GUI

This section explains how to create an EPG using the Cisco Cloud APIC GUI. Each service needs at least one consumer EPG and one provider EPG.

Before you begin

Create an application profile and a VRF.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Application Management**.

A list of **Application Management** options appear in the **Intent** menu.

Step 3 From the **Application Management** list in the **Intent** menu, click **Create EPG**. The **Create EPG** dialog box appears.

Step 4 Enter the appropriate values in each field as listed in the following *Create EPG Dialog Box Fields* table then continue.

Table 14: Create EPG Dialog Box Fields

Properties	Description
Name	Enter the name of the EPG.

Properties	Description
Tenant	To choose a tenant: <ol style="list-style-type: none"> a. Click Select Tenant. The Select Tenant dialog box appears. b. From the Select Tenant dialog, click to choose a tenant in the left column then click Select. You return to the Create EPG dialog box.
Application Profile	To choose an application profile: <ol style="list-style-type: none"> a. Click Select Application Profile. The Select Application Profile dialog box appears. b. From the Select Application Profile dialog, click to choose an application profile in the left column then click Select. You return to the Create EPG dialog box.
Description	Enter a description of the EPG.
Settings	
Type	Choose the EPG type: <ul style="list-style-type: none"> • Cloud - Click to create the EPG in the cloud. • External - Click to create an external EPG.
Route Reachability	(Visible when creating an external EPG) Click the Route Reachability drop-down list and choose: <ul style="list-style-type: none"> • On Premises • Internet • Unspecified
VRF	To choose a VRF: <ol style="list-style-type: none"> a. Click Select VRF. The Select VRF dialog box appears. b. From the Select VRF dialog, click to choose a VRF in the left column then click Select. You return to the Create EPG dialog box.

Properties	Description
Endpoint Selectors	

Properties	Description
	<p>Note See Configuring Instances in AWS, on page 82 for instructions on configuring instances in AWS as part of the endpoint selector configuration process.</p> <p>To add an endpoint selector:</p> <ol style="list-style-type: none"> a. Click Add Endpoint Selector to open the Add Endpoint Selector dialog. b. In the Add Endpoint Selector dialog, enter a name in the Name field. c. Click Selector Expression. The Key, Operator, and Value fields are enabled. d. Click the Key drop-down list to choose a key. The options are: <ul style="list-style-type: none"> • Choose IP if you want to use an IP address or subnet for the endpoint selector. • Choose Zone if you want to use an availability zone for the endpoint selector. • Choose Region if you want to use the Amazon Web Services region for the endpoint selector. • Choose Custom if you want to create a custom key for the endpoint selector. <p>Note When choosing the Custom option, the drop-down list becomes a text box. You need to enter a name for the key in the spaces after custom: (for example, custom: Location).</p> e. Click the Operator drop-down list to choose an operator. The options are: <ul style="list-style-type: none"> • equals: Used when you have a single value in the Value field. • not equals: Used when you have a single value in the Value field. • in: Used when you have multiple comma-separated values in the Value field. • not in: Used when you have multiple comma-separated values in the Value field. • has key: Used if the expression contains only a key.

Properties	Description
	<ul style="list-style-type: none"> • does not have key: Used if the expression contains only a key. <p>f. Enter a value in the Value field then click the check mark to validate the entries. The value you enter depends on the choices you made for the Key and Operator fields. For example, if the Key field is set to IP and the Operator field is set to equals, the Value field must be an IP address or subnet. However, if the Operator field is set to has key, the Value field is disabled.</p> <p>g. When finished, click the check mark to validate the selector expression.</p> <p>h. Determine if you want to create additional endpoint selector expressions to the endpoint selector. If you create more than one expression under a single endpoint selector, a logical AND exists between those expressions.</p> <p>For example, assume you created two sets of expressions under a single endpoint selector:</p> <ul style="list-style-type: none"> • Endpoint selector 1, expression 1: <ul style="list-style-type: none"> • Key: Zone • Operator: equals • Value: us-west-1a • Endpoint selector 1, expression 2: <ul style="list-style-type: none"> • Key: IP • Operator: equals • Value: 192.0.2.1/24 <p>In this case, if <i>both</i> of these expressions are true (if the availability zone is us-west-1a AND if the IP address belongs to subnet 192.0.2.1/24), then that endpoint is assigned to the Cloud EPG.</p>

Properties	Description
	<p>i. Click the check mark after every additional expression that you want to create under this endpoint selector then click Add when finished.</p> <p>If you create more than one endpoint selector under an EPG, a logical OR exists between those endpoint selectors. For example, assume you had created endpoint selector 1 as described in the previous step, and then you created a second endpoint selector as described below:</p> <ul style="list-style-type: none"> • Endpoint selector 2, expression 1: <ul style="list-style-type: none"> • Key: Region • Operator: in • Value: us-east-1, us-east-2 <p>In this case:</p> <ul style="list-style-type: none"> • If the availability zone is us-west-1a AND the IP address belongs to the 192.0.2.1/24 subnet (endpoint selector 1 expressions) <p>OR</p> <ul style="list-style-type: none"> • If the region is either us-east-1 or us-east-2 (endpoint selector 2 expression) <p>Then that end point is assigned to the Cloud EPG.</p>

Step 5 Click **Save** when finished.

Creating a Contract Using the Cisco Cloud APIC GUI

This section explains how to create a contract using the Cisco Cloud APIC GUI.

Before you begin

Create filters.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Application Management**.

A list of **Application Management** options appear in the **Intent** menu.

Step 3 From the **Application Management** list in the **Intent** menu, click **Create Contract**. The **Create Contract** dialog box appears.

Step 4 Enter the appropriate values in each field as listed in the following *Create Contract Dialog Box Fields* table then continue.

Table 15: Create Contract Dialog Box Fields

Properties	Description
Name	Enter the name of the contract.
Tenant	To choose a tenant: <ol style="list-style-type: none"> a. Click Select Tenant. The Select Tenant dialog box appears. b. From the Select Tenant dialog, click to choose a tenant in the left column then click Select. You return to the Create Contract dialog box.
Description	Enter a description of the contract.
Settings	
Scope	<p>The scope limits the contract to any endpoint groups within the same application profile, within the same VRF instance, throughout the fabric (globally), or within the same tenant.</p> <p>Note Shared services enables communication between EPGs in different tenants and between EPGs in different VRFs.</p> <p>To enable EPGs in one tenant to communicate with EPGs in another tenant, choose Global scope.</p> <p>To enable an EPG in one VRF to communicate with another EPG in a different VRF, choose Global or Tenant scope.</p> <p>For more information about shared services, see Shared Services, on page 44</p> <p>Click the drop-down arrow to choose from the following scope options:</p> <ul style="list-style-type: none"> • Application Profile • VRF • Global • Tenant
Apply Filter in Both Directions	<p>Put a check in the box to apply the same filters to traffic from consumer-to-provider and provider-to-consumer. Do not put a check in the box if you want to apply different filters for each direction of traffic.</p> <p>The check box is enabled by default.</p>

Properties	Description
Add Filter	<p>To choose a filter:</p> <ol style="list-style-type: none"> Click Add Filter. The filter row appears with a Select Filter option. Click Select Filter. The Select Filter dialog box appears. From the Select Filter dialog, click to choose a filter in the left column then click Select. You return to the Create Contract dialog box.

Step 5 Click **Save** when finished.

Specifying Consumer and Provider EPGs Using the Cisco Cloud APIC

This section explains how to specify an EPG as a consumer or a provider.

Before you begin

- You have configured a contract.
- You have configured an EPG.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Configuration**.

A list of **Configuration** options appears in the **Intent** menu.

Step 3 From the **Configuration** list in the **Intent** menu, click **EPG Communication**. The **EPG Communication** dialog box appears with the **Consumer EPGs**, **Contract**, and **Provider EPGs** information.

Step 4 To choose a contract:

- Click **Select Contract**. The **Select Contract** dialog appears.
- In the pane on the left side of the **Select Contract** dialog, click to choose a contract then click **Select**. The **Select Contract** dialog box closes.

Step 5 To add a consumer EPG:

- Click **Add Consumer EPGs**. The **Select Consumer EPGs** dialog appears.
- In the pane on the left side of the **Select Consumer EPGs** dialog, click to place a check in a check box to choose an EPG.

Step 6 To add a provider EPG:

- Click **Add Provider EPGs**. The **Select Provider EPGs** dialog appears.
- In the pane on the left side of the **Select Provider EPGs** dialog, click to place a check in a check box to choose a provider EPG.

- c) When finished, click **Select**. The **Select Provider EPGs** dialog box closes.

Creating a Filter Using the Cisco Cloud APIC GUI

This section explains how to create a filter using the Cisco Cloud APIC GUI.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Application Management**.

A list of **Application Management** options appear in the **Intent** menu.

Step 3 From the **Application Management** list in the **Intent** menu, click **Create Filter**. The **Create Filter** dialog box appears.

Step 4 Enter the appropriate values in each field as listed in the following *Create Filter Dialog Box Fields* table then continue.

Table 16: Create Filter Dialog Box Fields

Properties	Description
Name	Enter a name for the filter in the Name field.
Tenant	To choose a tenant: <ol style="list-style-type: none"> a. Click Select Tenant. The Select Tenant dialog box appears. b. From the Select Tenant dialog, click to choose a tenant in the left column then click Select. You return to the Create Filter dialog box.
Description	Enter a description of the filter.

Properties	Description
Add Filter	<p>To add a filter:</p> <ol style="list-style-type: none"> a. Click Add Filter Entry. The Create Filter Entry dialog box appears. b. Enter a name for the filter entry in the Name field. c. From the Select Filter dialog, click to choose a filter in the left column then click Select. You return to the Create Contract dialog box. d. Click the Ethernet Type drop-down list to choose an ethernet type. The options are: <ul style="list-style-type: none"> • IP • Unspecified <p>Note When Unspecified is chosen, the remaining fields are disabled.</p> e. Click the IP Protocol drop-down menu to choose a protocol. The options are: <ul style="list-style-type: none"> • icmp • tcp • udp • Unspecified <p>Note The remaining fields are enabled only when tcp or udp is chosen.</p> f. Enter the appropriate port information in the Origin Port from and to fields. g. Enter the appropriate port information in the Destination Port from and to fields. h. When finished entering filter entry information, click Add. You return to the Create Filter dialog box where you can repeat the steps to add another filter entry.

Step 5 When finished, click **Save**.

Creating a Cloud Context Profile Using the Cisco Cloud APIC GUI

This section explains how to create a cloud context profile using the Cisco Cloud APIC GUI.

Before you begin

Create a VRF.

- Step 1** Navigate to **Application Management > Cloud Context Profiles**.
The list of configure cloud context profiles appears.
- Step 2** Click **Actions > Create Cloud Context Profile**.
The **Create Cloud Context Profile** dialog box appears.
- Step 3** Enter the appropriate values in each field as listed in the following *Cloud Context Profile Dialog Box Fields* table then continue.

Table 17: Create Cloud Context Profile Dialog Box Fields

Properties	Description
Name	Enter the name of the cloud context profile.
Tenant	To choose a tenant: <ol style="list-style-type: none"> a. Click Select Tenant. The Select Tenant dialog box appears. b. From the Select Tenant dialog, click to choose a tenant in the left column then click Select. You return to the Create Cloud Context Profile dialog box.
Description	(Optional) Enter a description of the cloud context profile.
Settings	
Select Region	To choose a region: <ol style="list-style-type: none"> a. Click Select Region. The Select Region dialog box appears. b. From the Select Region dialog, click to choose a region in the left column then click Select. You return to the Create Cloud Context Profile dialog box.
Select VRF	To choose a VRF: <ol style="list-style-type: none"> a. Click Select VRF. The Select VRF dialog box appears. b. From the Select VRF dialog box, click to choose a VRF in the left column then click Select. You return to the Create Cloud Context Profile dialog box.

Properties	Description
Add CIDR	<p>Note The following subnets are reserved and should not be used in this Add CIDR field:</p> <ul style="list-style-type: none"> • 169.254.0.0/16 (reserved for VPN tunnel to the transit gateway) • 192.168.100.0/24 (reserved by the CCR for the bridge domain interface) <p>To add a CIDR:</p> <ol style="list-style-type: none"> a. Click Add CIDR. The Add CIDR dialog box appears. b. Enter the address in the Address field. c. Click Add Subnet and enter the subnet address in the Address field. d. To add availability zones: <ol style="list-style-type: none"> 1. Click Select Availability Zone. The Select Availability Zone dialog box appears. 2. From the Select Availability Zone dialog box, click to choose an availability zone in the left column. <p>Beginning with release 25.0(2), the type of availability zone shown in this window varies depending on the type of tenant that you selected for this cloud context profile.</p> <p>Note If you are creating a cloud context profile in a user tenant, you are restricted to only cloud availability zones in this window.</p> <p>See Availability Zones, on page 32 for more information.</p> <ol style="list-style-type: none"> 3. Click Select <p>You return to the Create Cloud Context Profile dialog box.</p> e. Click to check (enabled) or uncheck (disabled) the Primary check box. f. When finished, click Add.
VPN Gateway Router	(Optional) Click to check (enabled) or uncheck (disabled) in the VPN Gateway Router check box.
TGW Attachment	(Optional) Click to check (enabled) or uncheck (disabled) in the TGW Attachment check box.

Step 4 Click **Save** when finished.

Configuring Instances in AWS

When you configure endpoint selectors for Cisco Cloud APIC, you will also need to configure the instances that you will need in AWS that will correspond with the endpoint selectors that you configure for Cisco Cloud APIC.

This topic provides the instructions for configuring the instances in AWS. You can use these procedures to configure the instances in AWS either before you configure the endpoint selectors for Cisco Cloud APIC or afterward. For example, you might go to your account in AWS and create a custom tag or label in AWS first, then create an endpoint selector using a custom tag or label in Cisco Cloud APIC afterward. Or you might create an endpoint selector using a custom tag or label in Cisco Cloud APIC first, then go to your account in AWS and create a custom tag or label in AWS afterward.

Step 1

Review your cloud context profile configuration settings and determine which settings you will use with your AWS instance.

You must configure a cloud context profile as part of the AWS instance configuration process. When you configure a cloud context profile, the configurations, such as the VRF and region settings, are pushed out to AWS afterward.

- a) From the **Navigation** menu, choose the **Application Management** tab.

When the **Application Management** tab expands, a list of subtab options appear.

- b) Choose the **Cloud Context Profiles** subtab option.

A list of the cloud context profiles that you have created for your Cisco Cloud APIC are displayed.

- c) Select the cloud context profile that you will use as part of this AWS instance configuration process.

Various configuration parameters are displayed for this cloud context profile, such as the region, VRF, IP address and subnets. Use the information displayed in this window when you configure the AWS instance.

Step 2

Log in to the Amazon Web Services account for the Cisco Cloud APIC user tenant, if you are not logged in already.

Step 3

Go to **Services > EC2 > Instances > Launch Instance**.

Step 4

In the **Choose an Amazon Machine Image (AMI)** page, select an Amazon Machine Image (AMI).

Step 5

In the **Choose an Instance Type** page, select an instance type, then click **Configure Instance Details**.

Step 6

In the **Configure Instance Details** page, enter the necessary information in the appropriate fields.

- In the **Network** field, select your Cisco Cloud APIC VRF.

This would be the VRF that is associated with the cloud context profile that you are using as part of this AWS instance configuration process.

- In the **Subnet** field, select the subnet.

- In the **Auto-assign Public IP** field, if you want to have a public IP, select **Enable** from the scroll-down menu.

Step 7

When you have finished entering the necessary information into the **Configure Instance Details** page, click **Add Storage**.

Step 8

In the **Add Storage** page, accept the default values or configure the storage in this page, if necessary, and click **Add Tags**.

Step 9

In the **Add Tags** page, click **Add Tag** and enter the necessary information in the appropriate fields in this page.

Note If you will be using IP Address, Region or Zone for the type of endpoint selector later in these procedures, you do not have to enter any information in this page. In those situations, when you start the instance in AWS, the IP address, region or zone will be discovered by the Cisco Cloud APIC and the endpoint will be assigned to the EPG.

- **Key:** Enter the key that you will use when you create a custom tag for the type of endpoint selector that you are adding later in these procedures.

- **Value:** Enter the value that you will be using for this key.
- **Instances:** Check the box for this field.
- **Volumes:** Check the box for this field.

For example, if you are planning on creating a custom tag for a specific building for your endpoint selector later in these procedures (such as building6), you might enter the following values in these fields on this page:

- **Key:** Location
- **Value:** building6

Step 10 Click **Review and Launch**.

The **Select an existing key pair or create a new key pair** page appears. Use the information in this page if you want to ssh to the instance later on.

Creating a Backup Configuration Using the Cisco Cloud APIC GUI

This section explains how to create a backup configuration.

Before you begin

Create a remote location and a scheduler, if needed.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Operations**.

A list of **Operations** options appear in the **Intent** menu.

Step 3 From the **Operations** list in the **Intent** menu, click **Create Backup Configuration**. The **Create Backup Configuration** dialog box appears.

Step 4 Enter the appropriate values in each field as listed in the following *Create Backup Configuration Dialog Box Fields* table then continue.

Table 18: Create Backup Configuration Dialog Box Fields

Properties	Description
General	
Name	Enter the name of the backup configuration.
Description	Enter a description of the backup configuration.
Settings	
Backup Destination	Choose a backup destination. <ul style="list-style-type: none"> • Local • Remote

Properties	Description
Backup Object	

Properties	Description
	<p>Choose the root hierarchical content to consider for the backup</p> <ul style="list-style-type: none"> • Policy Universe • Selector Object—When chosen, this option adds the Object Type drop-down list and Object DN field. <ul style="list-style-type: none"> a. From the Object Type drop-down list, choose from the following options: <ul style="list-style-type: none"> • Tenant—When chosen the Select Tenant option appears. • Application Profile—When chosen the Select Application Profile option appears. • EPG—When chosen the Select EPG option appears. • Contract—When chosen the Select Contract option appears. • Filter—When chosen the Select Filter option appears. • VRF—When chosen the Select VRF option appears. • Device—When chosen the Select fvcloudLBCTX option appears. • Service Graph—When chosen the Select Service Graph option appears. • Cloud Context Profile—When chosen the Select Cloud Context Profile option appears. b. Click the Select <object_name>. The Select <object_name> dialog appears. c. From the Select <object_name> dialog, click to choose from the options in the left column then click Select. You return to the Create Backup Configuration dialog box. <p>Note The Object DN field is automatically populated with the DN of the object it will use as root of the object tree to backup</p> • Enter DN—When chosen, this option displays the Object DN field. <ul style="list-style-type: none"> a. From the Object DN field, enter the DN of a

Properties	Description
	specific object to use as the root of the object tree to backup.
Scheduler	<p>a. Click Select Scheduler to open the Select Scheduler dialog and choose a scheduler from the left-side column.</p> <p>b. Click the Select button at the bottom-right corner when finished.</p>
Trigger Backup After Creation	<p>Choose one of the following:</p> <ul style="list-style-type: none"> • Yes—(Default) Trigger a backup after creating the backup configuration. • No—Do not trigger a backup after creating the backup configuration.

Step 5 Click **Save** when finished.

Creating a Tech Support Policy Using the Cisco Cloud APIC GUI

This section explains how to create a tech support policy.

Before you begin

When creating a tech support policy for a remote location, you must first create the remote location.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Operations**.

A list of **Operations** options appear in the **Intent** menu.

Step 3 From the **Operations** list in the **Intent** menu, click **Create Tech Support**. The **Create Tech Support** dialog box appears.

Step 4 Enter the appropriate values in each field as listed in the following *Create Tech Support Dialog Box Fields* table then continue.

Table 19: Create Tech Support Dialog Box Fields

Properties	Description
General	
Name	Enter the name of the tech support policy.
Description	Enter a description of the tech support.
Settings	

Properties	Description
Export Destination	Choose an export destination. <ul style="list-style-type: none"> • Controller • Remote Location—When chosen the Select Remote Location option appears. <ol style="list-style-type: none"> Click Select Remote Location. The Select Remote Location dialog box appears. From the Select Remote Location dialog, click to choose a remote location in the left column then click Select. You return to the Create Tech Support dialog box.
Include Pre-Upgrade Logs	Click to place a check in the Enabled check box if you want to include pre-upgrade logs in the tech support policy.
Trigger After Creation	Click to place a check in the Enabled (the default) check box if you want to create the tech support policy after the policy creation. To disable, click the check box to uncheck.

Step 5 Click **Save** when finished.

Creating a Trigger Scheduler Using the Cisco Cloud APIC GUI

This section explains how to create a trigger scheduler.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Operations**.

A list of **Operations** options appear in the **Intent** menu.

Step 3 From the **Operations** list in the **Intent** menu, click **Create Scheduler**. The **Create Trigger Scheduler** dialog box appears.

Step 4 Enter the appropriate values in each field as listed in the following *Create Trigger Scheduler Dialog Box Fields* table then continue.

Table 20: Create Trigger Scheduler Dialog Box Fields

Properties	Description
General	
Name	Enter the name of the trigger scheduler policy.
Description	Enter a description of the trigger scheduler.
Settings	

Properties	Description
Recurring Windows	<p>Click Add Recurring Window. The Add Recurring Window dialog appears.</p> <ol style="list-style-type: none"> a. From the Schedule drop-down list, choose from the following. <ul style="list-style-type: none"> • every-day • Monday • Tuesday • Wednesday • Thursday • Friday • Saturday • Sunday • odd-day • even-day b. From the Start Time field, enter a time. c. From the Maximum Concurrent Tasks field, enter a number or leave the field empty to specify unlimited. d. From the Maximum Running Time, click to choose Unlimited or Custom. e. Click Add when finished.
Add One Time Window	<p>Click Add One Time Window. The Add One Time Window dialog appears.</p> <ol style="list-style-type: none"> a. From the Start Time field, enter a date and time. b. From the Maximum Concurrent Tasks field, enter a number or leave the field blank to specify unlimited. c. From the Maximum Running Time, click to choose Unlimited or Custom. d. Click Add when finished.

Step 5 Click **Save** when finished.

Creating a Remote Location Using the Cisco Cloud APIC GUI

This section explains how to create a remote location using the Cisco Cloud APIC.

- Step 1** Click the **Intent** icon. The **Intent** menu appears.
- Step 2** Click the drop-down arrow below the **Intent** search box and choose **Operations**.
A list of **Operations** options appear in the **Intent** menu.
- Step 3** From the **Operations** list in the **Intent** menu, click **Create Remote Location**. The **Create Remote Location** dialog box appears.
- Step 4** Enter the appropriate values in each field as listed in the following *Create Remote Location Dialog Box Fields* table then continue.

Table 21: Create Remote Location Dialog Box Fields

Properties	Description
General	
Name	Enter the name of the remote location policy.
Description	Enter a description of the remote location policy.
Settings	
Hostname/IP Address	Enter the hostname or IP address of the remote location
Protocol	Choose a protocol: <ul style="list-style-type: none"> • FTP • SFTP • SCP
Path	Enter the path for the remote location.
Port	Enter the port for the remote location.
Username	Enter a username for the remote location.
Authentication Type	When using SFTP or SCP, choose the authentication type: <ul style="list-style-type: none"> • Password • SSH Key
SSH Key Content	Enter the SSH key content.
SSH Key Passphrase	SSH key passphrase.
Password	Enter a password for accessing the remote location.
Confirm Password	Reenter the password for accessing the remote location.

Properties	Description
Management EPG	<ol style="list-style-type: none"> a. Click Select Management EPG. The Select Management EPG dialog appears. b. From the column on the left, click to choose a management EPG. c. Click Select.

Step 5 Click **Save** when finished.

Creating a Login Domain Using the Cisco Cloud APIC GUI

This section explains how to create a login domain using the Cisco Cloud APIC GUI.

Before you begin

Create a provider before creating a non-local domain.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Administrative**.

A list of **Administrative** options appear in the **Intent** menu.

Step 3 From the **Administrative** list in the **Intent** menu, click **Create Login Domain**. The **Create Login Domain** dialog box appears.

Step 4 Enter the appropriate values in each field as listed in the following *Create Login Domain Dialog Box Fields* table then continue.

Table 22: Create Login Domain Dialog Box Fields

Properties	Description
Name	Enter the name of the login domain.
Description	Enter a description of the login domain.
Realm	Choose a realm: <ul style="list-style-type: none"> • Local • LDAP—Requires adding providers and choosing an authentication type. • RADIUS—Requires adding providers. • TACACS+—Requires adding providers. • SAML—Requires adding providers.

Properties	Description
Providers	<p>To add a provider:</p> <ol style="list-style-type: none"> a. Click Add Providers. The Select Providers dialog appears with a list of providers in the left pane. b. Click to choose a provider. c. Click Select to add the provider.
Advanced Settings	Displays the Authentication Type and LDAP Group Map Rules fields.
Authentication Type	<p>When LDAP is chosen for realm option, choose one of the following authentication types:</p> <ul style="list-style-type: none"> • Cisco AV Pairs—(Default) • LDAP Group Map Rules—Requires adding LDAP group map rules.

Properties	Description
<p>LDAP Group Map Rules</p>	<p>To add an LDAP group map rule:</p> <ol style="list-style-type: none"> a. Click Add LDAP Group Map Rule. The Add LDAP Group Map Rule dialog appears with a list of providers in the left pane. b. Enter a name for the rule in the Name field. c. Enter a description for the rule in the Description field. d. Enter a group DN for the rule in the Group DN field. e. Add security domains: <ol style="list-style-type: none"> 1. Click Add Security Domain. The Add Security Domain dialog box appears. 2. Click Select Security Domain. The Select Security Domain dialog box appears with a list of security domains in the left pane. 3. Click to choose a security domain. 4. Click Select to add the security domain. You return to the Add Security Domain dialog box. 5. Add a user role: <ol style="list-style-type: none"> a. From the Add Security Domain dialog box, click Select Role. The Select Role dialog box appears with a list of roles in the left pane. b. Click to choose a role. c. Click Select to add the role. You return to the Add Security Domain dialog box. d. From the Add Security Domain dialog box, click the Privilege Type drop-down list and choose Read Privilege or Write Privilege. e. Click the check mark on the right side of the Privilege Type drop-down list to confirm. f. Click Add when finished. You return to the Add LDAP Group Map Rule dialog box where you can add another security domain.

Step 5 Click **Save** when finished.

Creating a Provider Using the Cisco Cloud APIC GUI

This section explains how to create a provider using the Cisco Cloud APIC GUI.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Administrative**.

A list of **Administrative** options appear in the **Intent** menu.

Step 3 From the **Administrative** list in the **Intent** menu, click **Create Provider**. The **Create Provider** dialog box appears.

Step 4 Enter the appropriate values in each field as listed in the following *Create Provider Dialog Box Fields* table then continue.

Table 23: Create Provider Dialog Box Fields

Properties	Description
Hostname/IP Address	Enter the hostname or IP address of the provider.
Description	Enter a description of the provider.
Type	Click the Type drop-down list and choose one of the following types: <ul style="list-style-type: none"> • LDAP • RADIUS • TACACS+ • SAML <p>Note A set of fields will appear based on the type that you choose.</p>
[LDAP] Settings	
Bind DN	Enter the LDAP bind DN.
Base DN	Enter the LDAP base DN.
Password	Enter a password for the LDAP settings.
Confirm Password	Reenter the password for the LDAP settings.
Port	Enter the port number for the provider type.
Advanced Settings	Displays additional fields in the Settings section of the provider dialog box.
Timeout (sec)	Enter the number of seconds allowed before a timeout occurs. The default is 30.
Retries	Enter the number of allowed retries. The default is 1.

Properties	Description
SSL	To enable SSL, click to place a check in the SSL check box. To disable SSL, click to remove the check from the SSL check box. The default is enabled.
SSL Certificate Validation Level	Choose one of the following: <ul style="list-style-type: none"> • Permissive • Strict
Attribute	Enter an LDAP attribute in the Attribute text box.
Filter Type	Choose a filter type: <ul style="list-style-type: none"> • Default • Microsoft AD • Custom
Filter	Enter an LDAP filter in the text box. This option only appears when the Custom filter type is chosen.
Select Management EPG	To add a management EPG: <ol style="list-style-type: none"> Click Select Management EPG. The Select Management EPG dialog appears with a list of EPGs in the left pane. Click to choose an EPG. Click Select to add the management EPG to the LDAP.
Server Monitoring	To enable server monitoring, click to place a check in the Enabled check box. To disable server monitoring, click to remove the check from the Enabled check box. The default is disabled.
[RADIUS] Settings	
Key	Enter the RADIUS key.
Confirm Key	Reenter the RADIUS key.
Advanced Settings	Displays additional fields in the Settings section of the provider dialog box.
Port	Enter the port number for the RADIUS settings. The default is 1812.

Properties	Description
Authentication Protocol	Choose from the following: <ul style="list-style-type: none"> • PAP—(Default) • CHAP • MS-CHAP
Timeout (sec)	Enter the number of seconds allowed before a timeout occurs. The default is 5.
Retries	Enter the number of allowed retries. The default is 1.
Select Management EPG	To add a management EPG: <ol style="list-style-type: none"> Click Select Management EPG. The Select Management EPG dialog appears with a list of EPGs in the left pane. Click to choose an EPG. Click Select to add the management EPG to the RADIUS.
Server Monitoring	To enable server monitoring, click to place a check in the Enabled check box. To disable server monitoring, click to remove the check from the Enabled check box. The default is disabled.
[TACACS+] Settings	
Key	Enter the TACACS+ key.
Confirm Key	Reenter the TACACS+ key.
Advanced Settings	Displays additional fields in the Settings section of the provider dialog box.
Port	Enter the port number for the TACACS+ settings. The default is 1812.
Authentication Protocol	Choose from the following: <ul style="list-style-type: none"> • CHAP • MS-CHAP • PAP—(Default)
Timeout (sec)	Enter the number of seconds allowed before a timeout occurs. The default is 5.
Retries	Enter the number of allowed retries. The default is 1.

Properties	Description
Select Management EPG	To add a management EPG: <ol style="list-style-type: none"> Click Select Management EPG. The Select Management EPG dialog appears with a list of EPGs in the left pane. Click to choose an EPG. Click Select to add the management EPG to the TACACS+.
Server Monitoring	To enable server monitoring, click to place a check in the Enabled check box. To disable server monitoring, click to remove the check from the Enabled check box. The default is disabled.
[SAML] Settings	
Identity Provider	Choose from the following identity providers: <ul style="list-style-type: none"> • ADFS—(default) • OKTA • PING IDENTITY
Identity Provider Metadata URL	Enter the metadata URL provided by the identity provider.
Entity ID	Enter a unique ID as the SAML entity identifier.
HTTPS Proxy for Metadata URL	Enter the HTTPS proxy used to reach the identity provider's metadata URL.
Advanced Settings	Displays additional fields in the Settings section of the provider dialog box.
GUI Redirect Banner Message (URL)	Enter the GUI redirect banner message.
Certificate Authority	To choose a certificate authority: <ol style="list-style-type: none"> Click Select Certificate Authority. The Select Certificate Authority dialog appears with a list of certificates in the left pane. Click to choose a certificate. Click Select to add the certificate. You return to the Create Provider dialog box.
Timeout (sec)	Enter the number of seconds allowed before a timeout occurs. The default is 5.
Retries	Enter the number of allowed retries. The default is 1.

Properties	Description
Signature Algorithm Authentication User Requests*	Click the Signature Algorithm for Requests drop-down list and choose one of the following: <ul style="list-style-type: none"> • RSA SHA1 • RSA SHA224 • RSA SHA256 (Default) • RSA SHA384 • RSA SHA512
Sign SAML Authentication Requests	To enable, click to place a check in the check box. To disable, click to remove the check from the check box. The default is enabled.
Sign SAML Response Message	To enable, click to place a check in the check box. To disable, click to remove the check from the check box. The default is enabled.
Sign Assertions in SAML Response	To enable, click to place a check in the check box. To disable, click to remove the check from the check box. The default is enabled.
Encrypt SAML Assertions	To enable, click to place a check in the check box. To disable, click to remove the check from the check box. The default is enabled.

Step 5 Click **Save** when finished.

Creating a Security Domain Using the Cisco Cloud APIC GUI

A security domain restricts the tenant to the security domains that you add. If you do not add a security domain, all security domains will have access to this tenant. This section explains how to create a security domain using the GUI.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Administrative**.

A list of **Administrative** options appear in the **Intent** menu.

Step 3 From the **Administrative** list in the **Intent** menu, click **Create Security Domain**. The **Create Security Domain** dialog box appears.

Step 4 In the **Name** field, enter the name of the security domain.

Step 5 In the **Description** field, enter a description of the security domain.

Step 6 Click **Save** when finished.

Creating a Role Using the Cisco Cloud APIC GUI

This section explains how to create a role using the Cisco Cloud APIC GUI.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Administrative**.

A list of **Administrative** options appear in the **Intent** menu.

Step 3 From the **Administrative** list in the **Intent** menu, click **Create Role**. The **Create Role** dialog box appears.

Step 4 Enter the appropriate values in each field as listed in the following *Create Role Dialog Box Fields* table then continue.

Table 24: Create Role Dialog Box Fields

Properties	Description
General	
Name	Enter a name for the role in the Name field.
Description	Enter a description of the role.
Settings	

Properties	Description
Privilege	

Properties	Description
	<p>Click to place a check mark in the check boxes of the privileges you want to assign the user. The privileges are:</p> <ul style="list-style-type: none"> • aaa—Used for configuring authentication, authorization, accounting and import/export policies. • access-connectivity-11—Used for Layer 1 configuration under infra. Example: selectors and port Layer 1 policy configurations. • access-connectivity-12—Used for Layer 2 configuration under infra. Example: Encap configurations on selectors, and attachable entity. • access-connectivity-13—Used for Layer 3 configuration under infra and static route configurations under a tenant's L3Out. • access-connectivity-mgmt—Used for management infra policies. • access-connectivity-util—Used for tenant ERSPAN policies. • access-equipment—Used for access port configuration. • access-protocol-11—Used for Layer 1 protocol configurations under infra. • access-protocol-12—Used for Layer 2 protocol configurations under infra. • access-protocol-13—Used for Layer 3 protocol configurations under infra. • access-protocol-mgmt—Used for fabric-wide policies for NTP, SNMP, DNS, and image management. • access-protocol-ops—Used for operations-related access policies such as cluster policy and firmware policies. • access-protocol-util—Used for tenant ERSPAN policies. • access-qos—Used for changing CoPP and QoS-related policies. • admin—Complete access to everything (combine ALL roles) • fabric-connectivity-11—Used for Layer 1 configuration under the fabric. Example: selectors and port Layer 1 policy and vPC protection.

Properties	Description
	<ul style="list-style-type: none"> • fabric-connectivity-l2—Used in firmware and deployment policies for raising warnings for estimating policy deployment impact. • fabric-connectivity-l3—Used for Layer 3 configuration under the fabric. Example: Fabric IPv4 and MAC protection groups. • fabric-connectivity-mgmt—Used for atomic counter and diagnostic policies on leaf switches and spine switches. • fabric-connectivity-util—Used for atomic counter, diagnostic, and image management policies on leaf switches and spine switches. • fabric-equipment—Used for atomic counter, diagnostic, and image management policies on leaf switches and spine switches. • fabric-protocol-l1—Used for Layer 1 protocol configurations under the fabric. • fabric-protocol-l2—Used for Layer 2 protocol configurations under the fabric. • fabric-protocol-l3—Used for Layer 3 protocol configurations under the fabric. • fabric-protocol-mgmt—Used for fabric-wide policies for NTP, SNMP, DNS, and image management. • fabric-protocol-ops—Used for ERSPAN and health score policies. • fabric-protocol-util—Used for firmware management traceroute and endpoint tracking policies. • none—No privilege. • nw-svc-device—Used for managing Layer 4 to Layer 7 service devices. • nw-svc-devshare—Used for managing shared Layer 4 to Layer 7 service devices. • nw-svc-params—Used for managing Layer 4 to Layer 7 service policies. • nw-svc-policy—Used for managing Layer 4 to Layer 7 network service orchestration.

Properties	Description
	<ul style="list-style-type: none"> • ops—Used for operational policies including monitoring and troubleshooting policies such as atomic counter, SPAN, TSW, tech support, traceroute, analytics, and core policies. • tenant-connectivity-l1—Used for Layer 1 connectivity changes, including bridge domains and subnets. • tenant-connectivity-l2—Used for Layer 2 connectivity changes, including bridge domains and subnets. • tenant-connectivity-l3—Used for Layer 3 connectivity changes, including VRFs. • tenant-connectivity-mgmt—Used for tenant in-band and out-of-band management connectivity configurations and for debugging/monitoring policies such as atomic counters and health score. • tenant-connectivity-util—Used for atomic counter, diagnostic, and image management policies on leaf switches and spine switches. • tenant-epg—Used for managing tenant configurations such as deleting/creating endpoint groups, VRFs, and bridge domains. • tenant-ext-connectivity-l2—Used for managing tenant L2Out configurations. • tenant-ext-connectivity-l3—Used for managing tenant L3Out configurations. • tenant-ext-connectivity-mgmt—Used as write access for firmware policies. • tenant-ext-connectivity-util—Used for debugging/monitoring/observer policies such as traceroute, ping, oam, and eptrk. • tenant-ext-protocol-l1—Used for managing tenant external Layer 1 protocols. Generally only used for write access for firmware policies. • tenant-ext-protocol-l2—Used for managing tenant external Layer 2 protocols. Generally only used for write access for firmware policies. • tenant-ext-protocol-l3—Used for managing tenant external Layer 3 protocols such as BGP, OSPF, PIM, and IGMP. • tenant-ext-protocol-mgmt—Used as write access for firmware policies.

Properties	Description
	<ul style="list-style-type: none"> • tenant-ext-protocol-util—Used for debugging/monitoring/observer policies such as traceroute, ping, oam, and eptrk. • tenant-network-profile—Used for managing tenant configurations, such as deleting and creating network profiles, and deleting and creating endpoint groups. • tenant-protocol-11—Used for managing configurations for Layer 1 protocols under a tenant. • tenant-protocol-12—Used for managing configurations for Layer 2 protocols under a tenant. • tenant-protocol-13—Used for managing configurations for Layer 3 protocols under a tenant. • tenant-protocol-mgmt—Only used as write access for firmware policies. • tenant-protocol-ops—Used for tenant traceroute policies. • tenant-protocol-util—Used for debugging/monitoring/observer policies such as traceroute, ping, oam, and eptrk. • tenant-qos—Only used as Write access for firmware policies. • tenant-security—Used for Contract related configurations for a tenant. • vmm-connectivity—Used to read all the objects in APIC's VMM inventory required for VM connectivity. • vmm-ep—Used to read VM and Hypervisor endpoints in the APIC's VMM inventory. • vmm-policy—Used for managing policies for VM networking. • vmm-protocol-ops—Not used by VMM policies. • vmm-security—Used for Contract related configurations for a tenant.

Step 5 Click **Save** when finished.

Creating an RBAC Rule Using the Cisco Cloud APIC GUI

This section explains how to create an RBAC rule using the GUI.

Before you begin

Create a security domain.

-
- Step 1** Click the **Intent** icon. The **Intent** menu appears.
- Step 2** Click the drop-down arrow below the **Intent** search box and choose **Administrative**.
A list of **Administrative** options appears in the **Intent** menu.
- Step 3** From the **Administrative** list in the **Intent** menu, click **Create RBAC Rule**. The **Create RBAC Rule** dialog box appears.
- Step 4** In the **DN** field, enter the DN for the rule.
- Step 5** Choose a security domain:
- Click **Select Security Domain**. The **Select Security Domain** dialog box appears.
 - From the **Select Security Domain** dialog box, click to choose a security domain from the column on the left then click **Select**. You return to the **Create RBAC Rule** dialog box.
- Step 6** From the **Allow Writes** field, click **Yes** to allow writes or **No** to not allow writes.
- Step 7** Click **Save** when finished.
-

Creating a Certificate Authority Using the Cisco Cloud APIC GUI

This section explains how to create a certificate authority using the GUI.

Before you begin

- Have the certificate chain.
- If the certificate authority is for a tenant, create the tenant.

-
- Step 1** Click the **Intent** icon. The **Intent** menu appears.
- Step 2** Click the drop-down arrow below the **Intent** search box and choose **Administrative**.
A list of **Administrative** options appears in the **Intent** menu.
- Step 3** From the **Administrative** list in the **Intent** menu, click **Create Certificate Authority**. The **Create Certificate Authority** dialog box appears.
- Step 4** Enter the appropriate values in each field as listed in the following *Create Certificate Authority Dialog Box Fields* table then continue.

Table 25: Create Certificate Authority Dialog Box Fields

Properties	Description
Name	Enter the name of the certificate authority.
Description	Enter a description of the certificate authority.

Properties	Description
Used for	Choose from the following options: <ul style="list-style-type: none"> • Tenant—Choose if the certificate authority is for a specific tenant. When chosen, the Select Tenant option appears in the GUI. • System—Choose if the certificate authority is for the system.
Select Tenant	To choose a tenant: <ol style="list-style-type: none"> Click Select Tenant. The Select Tenant dialog box appears. From the Select Tenant dialog, click to choose a tenant in the left column then click Select. You return to the Create Certificate Authority dialog box.
Certificate Chain	Enter the certificate chain in the Certificate Chain text box. <p>Note Add the certificates for a chain in the following order:</p> <ol style="list-style-type: none"> CA Sub-CA Subsub-CA Server

Step 5 Click **Save** when finished.

Creating a Key Ring Using the Cisco Cloud APIC GUI

This section explains how to create a key ring using the Cisco Cloud APIC GUI.

Before you begin

- Create a certificate authority.
- Have a certificate.
- If the key ring is for a specific tenant, create the tenant.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Administrative**.

A list of **Administrative** options appear in the **Intent** menu.

- Step 3** From the **Administrative** list in the **Intent** menu, click **Create Key Ring**. The **Create Key Ring** dialog box appears.
- Step 4** Enter the appropriate values in each field as listed in the following *Create Key Ring Dialog Box Fields* table then continue.

Table 26: Create Key Ring Dialog Box Fields

Properties	Description
Name	Enter the name of the key ring.
Description	Enter a description of the key ring.
Used for	<ul style="list-style-type: none"> • System—The key ring is for the system. • Tenant—The key ring is for a specific tenant. Displays a Tenant field for specifying the tenant.
Select Tenant	<p>To choose a tenant:</p> <ol style="list-style-type: none"> a. Click Select Tenant. The Select Tenant dialog box appears. b. From the Select Tenant dialog, click to choose a tenant in the left column then click Select. You return to the Create Key Ring dialog box.
Settings	
Certificate Authority	<p>To choose a certificate authority:</p> <ol style="list-style-type: none"> a. Click Select Certificate Authority. The Select Certificate Authority dialog appears. b. Click to choose a certificate authority in the column on the left. c. Click Select. You return to the Create Key Ring dialog box.
Private Key	<p>Choose one of the following:</p> <ul style="list-style-type: none"> • Generate New Key—Generates a new key. • Import Existing Key—Displays the Private Key text box and enables you to use an existing key.
Private Key	Enter an existing key in the Private Key text box (for the Import Existing Key option).

Properties	Description
Modulus	Click the Modulus drop-down list to choose from the following: <ul style="list-style-type: none"> • MOD 512 • MOD 1024 • MOD 1536 • MOD 2048—(Default)
Certificate	Enter the certificate information in the Certificate text box.

Step 5 Click **Save** when finished.

Creating a Local User Using the Cisco Cloud APIC GUI

This section explains how to create a local user using the Cisco Cloud APIC GUI.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Administrative**.

A list of **Administrative** options appear in the **Intent** menu.

Step 3 From the **Administrative** list in the **Intent** menu, click **Create Local User**. The **Create Local User** dialog box appears.

Step 4 Enter the appropriate values in each field as listed in the following *Create Local User Dialog Box Fields* table then continue.

Table 27: Create Local User Dialog Box Fields

Properties	Description
Name	Enter the username of the local user.
Password	Enter the password for the local user.
Confirm Password	Reenter the password for the local user.
Description	Enter a description of the local user.
Settings	
Account Status	To choose the account status: <ul style="list-style-type: none"> • Active—Activates the local user account. • Inactive—Deactivates the local user account.
First Name	Enter the first name of the local user.

Properties	Description
Last Name	Enter the last name of the local user.
Email Address	Enter the email address of the local user.
Phone Number	Enter the phone number of the local user.
Security Domains	<p>To add a security domain:</p> <ol style="list-style-type: none"> a. Click Add Security Domain. The Add Security Domain dialog box appears. b. Click Select Security Domain. The Select Security Domain dialog box appears with a list of security domains in the left pane. c. Click to choose a security domain. d. Click Select to add the security domain. You return to the Add Security Domain dialog box. e. Add a user role: <ol style="list-style-type: none"> 1. From the Add Security Domain dialog box, click Select Role. The Select Role dialog box appears with a list of roles in the left pane. 2. Click to choose a role. 3. Click Select to add the the role. You return to the Add Security Domain dialog box. 4. From the Add Security Domain dialog box, click the Privilege Type drop-down list and choose Read Privilege or Write Privilege. 5. Click the check mark on the right side of the Privilege Type drop-down list to confirm. 6. Click Add when finished. You return to the Create Local User dialog box where you can add another security domain.

Step 5 Click **Advanced Settings** and enter the appropriate values in each field as listed in the following *Create Local User Dialog Box Fields: Advanced Settings* table then continue.

Table 28: Create Local User Dialog Box Fields: Advanced Settings

Property	Description
Account Expires	If you choose Yes , the account is set to expire at the time that you choose.
Password Update Required	If you choose Yes , the user must change the password upon the next login.

Property	Description
OTP	Put a check in the box to enable the one-time password feature for the user.
User Certificates	To add a user certificate: <ol style="list-style-type: none"> a. Click Add X509 Certificate. The Add X509 Certificate dialog box appears. b. Enter a name in the Name field. c. Enter the X509 certificate in the User X509 Certificate text box. d. Click Add. The X509 certificate in the User X509 Certificate dialog box closes. You return to the Local User dialog box.
SSH Keys	To add a an SSH key: <ol style="list-style-type: none"> a. Click Add SSH Key. The Add SSH Key dialog box appears. b. Enter a name in the Name field. c. Enter the SSH key in the Key text box. d. Click Add. The Add SSH Key dialog box closes. You return to the Local User dialog box.

Step 6 Click **Save** when finished.

Managing Regions (Configuring a Cloud Template) Using the Cisco Cloud APIC GUI

Regions are configured during the first-time setup. When configured, you specify the regions that are managed by Cisco Cloud APIC and the region's inter-site and inter-region connectivity. This section explains how to manage regions with the cloud template using the Cisco Cloud APIC GUI after the initial installation.

For more information about cloud templates, see [About the Cloud Template, on page 39](#).

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Configuration**.

A list of options appear in the **Intent** menu.

Step 3 From the **Configuration** list in the **Intent** menu, click **Cloud APIC Setup**. The **Set up - Overview** dialog box appears.

Step 4 In the **Region Management** area, click **Edit Configuration**.

The **Setup - Region Management** dialog box appears, and the first step in the **Setup - Region Management** series of steps appears, **Regions to Manage**, with a list of managed regions.

Step 5 If you want inter-site connectivity, click to place a check mark in the **Enabled** box in the **Inter-Site Connectivity** area. The **Inter-Site Connectivity** step is added in the **Setup - Region Management** steps at the top of the page.

Step 6 To choose a region that you want to be managed by the Cisco Cloud APIC, click to place a check mark in check box of that region.

Step 7 To deploy cloud routers locally to this region, click to place a check mark in the **Cloud Routers** check box for that region.

Step 8 To configure the fabric infra connectivity for the cloud site, click **Next**.
The next step in the **Setup - Region Management** series of steps appears, **General Connectivity**

Step 9 To add a subnet pool for the CCRs, click **Add Subnet Pool for Cloud Router** and enter the subnet in the text box.

Note The /24 subnet provided during the Cloud APIC deployment would be sufficient for up to two cloud sites. If you need to manage more than two cloud sites, you need to add more subnets.

Step 10 Enter a value in the **BGP Autonomous System Number for CCRs** field.

The BGP ASN can be in the range of 1 - 65534.

Step 11 In the **Assign Public IP to CCR Interface** field, determine if you want to have a public or a private IP address assigned to the CCR interface.

- To have a public IP address assigned to the CCR interface, leave the check in the **Enabled** check box. By default, the **Enabled** check box is checked.
- To have public IP disabled to the CCR interfaces, uncheck the **Enabled** check box. A private IP address is used for connectivity in this case.

Note Disabling or enabling a public IP address is a disruptive operation and can result in traffic loss.

Beginning with release 5.2(1), both the public and private IP addresses assigned to a CCR are displayed with the other details of the router in the Cloud Resources area. If a public IP is not assigned to a CCR, only the private IP is displayed.

Step 12 To choose the number of routers per region, click the **Number of Routers Per Region** drop-down list and click **2**, **3**, or **4**.

Step 13 Enter a username in the **Username** text box.

Step 14 Enter a password in the **Password** and **Confirm Password** text boxes.

Step 15 To choose the throughput value, click the **Throughput of the routers** drop-down list.

Note

- Cloud routers should be undeployed from all regions before changing the throughput or login credentials.
- Beginning with release 25.0(3), Cisco Cloud APIC moves from the Cisco Cloud Services Router 1000v to the Cisco Catalyst 8000V. For information on the throughput values for the Cisco Catalyst 8000V, see [About the Cisco Catalyst 8000V, on page 24](#).

Step 16 (Optional) To specify the license token, enter the product instance registration token in the **License Token** text box.

- Note**
- Beginning with release 25.0(3), Cisco Cloud APIC moves from the Cisco Cloud Services Router 1000v to the Cisco Catalyst 8000V. For licensing information for the Cisco Catalyst 8000V, see [About the Cisco Catalyst 8000V, on page 24](#).
 - If no token is entered, the CCR will be in EVAL mode.
 - If the public IP addresses are disabled to the CCRs in [Step 11, on page 111](#), the only supported option is **AWS Direct Connect or Azure Express Route to Cisco Smart Software Manager (CSSM)** when registering smart licensing for CCRs with private IP addresses (available by navigating to **Administrative > Smart Licensing**). You must provide reachability to the CSSM through AWS Direct Connect or Azure Express Route in this case. When the public IP addresses are disabled, public internet cannot be used because private IP addresses are being used. The connectivity should therefore use Private Connection, which is AWS Direct Connect or Azure Express Route.

Step 17 Click **Next**.

- If you placed a check mark in the **Enabled** box in the **Inter-Site Connectivity** area earlier in these procedures, **Inter-Site Connectivity** appears as the next step in the **Setup - Region Management** series of steps. Go to [Step 18, on page 112](#).
- If you did not place a check mark in the **Enabled** box in the **Inter-Site Connectivity** area earlier in these procedures, go to [Step 22, on page 112](#).

Step 18 To enter a peer public IP address of the IPsec Tunnel peer on-premises in the text box, click **Add Public IP of IPsec Tunnel Peer**.

Step 19 Enter the OSPF area ID in the **OSPF Area Id** text box.

Step 20 To add an external subnet pool, click **Add External Subnet** and enter a subnet pool in the text box.

Step 21 When you have configured all the connectivity options, click **Next** at the bottom of the page.

Step 22 Click **Save and Continue** when finished.

Configuring Cisco Cloud APIC Using the REST API

Creating a Tenant Using the REST API

This section demonstrates how to create a tenant and assigns using the REST API.

To create a tenant:

```
<polUni>
  <fvTenant name="infra">
    <cloudAwsProvider region="us-east-1" accessKeyId="123" secretAccessKey="ABCDE" providerId="admin"
    status=""/>
  </fvTenant>
</polUni>
```

Creating a Contract Using the REST API

This example demonstrates how to create a contract for the Cisco Cloud APIC using the REST API.

Before you begin

Create filters.

To create a contract:

Example:

```
<polUni>
  <fvTenant name="t2" status="">
    <vzFilter descr="" name="http-family-destination" ownerKey="" ownerTag="">
      <vzEntry name="http" prot="tcp" etherT="ip" dFromPort="http" dToPort="http"/>
      <vzEntry name="https" prot="tcp" etherT="ip" dFromPort="https" dToPort="https"/>
    </vzFilter>
    <vzBrCP name="httpFamily">
      <vzSubj name="default" revFltPorts="yes" targetDscp="unspecified">
        <vzRsSubjFiltAtt action="permit" directives="" tnVzFilterName="http-family-destination"/>
      </vzSubj>
    </vzBrCP>
  </fvTenant>
</polUni>
```

Creating a Cloud Context Profile Using the REST API

This section demonstrates how to create a cloud context profile.

Before you begin

Create a VRF.

Step 1 For releases prior to release 25.0(2), enter a post similar to the following to create a cloud context profile:

Example:

```
<polUni>
<fvTenant name="Corp1" status="">
  <cloudAwsProvider accessKeyId="" secretAccessKey="" providerId="aws" status="" accountId="">

    <fvCtx name="prod-1" status="">
      <bgpRtTargetP af="ipv4-ucast">
        <bgpRtTarget rt="route-target:as4-nn2:400:400" type="export"/>
        <bgpRtTarget rt="route-target:as4-nn2:400:400" type="import"/>
      </bgpRtTargetP>
    </fvCtx>

    <fvCtx name="prod-2" status="">
      <bgpRtTargetP af="ipv4-ucast">
        <bgpRtTarget rt="route-target:as4-nn2:500:500" type="export"/>
        <bgpRtTarget rt="route-target:as4-nn2:500:500" type="import"/>
      </bgpRtTargetP>
    </fvCtx>
  </cloudAwsProvider>
</fvTenant>
```

```

<cloudVpnGwPol name="VgwPol" status=""/>

<cloudApp name="payment" status="">
  <cloudEPg name="web" status="">
    <cloudRsCloudEPgCtx tnFvCtxName="prod-1" />
  </cloudEPg>
</cloudApp>
<cloudApp name="billing">
  <cloudEPg name="app">
    <cloudRsCloudEPgCtx tnFvCtxName="prod-2" />
  </cloudEPg>
</cloudApp>

<cloudCtxProfile name="prod-web-east-1">
  <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-aws/region-us-east-1"/>
  <cloudRsToCtx tnFvCtxName="prod-1"/>
  <cloudRouterP name="RouterP1" type="vpn-gw">
    <cloudRsToVpnGwPol tnCloudVpnGwPolName="VgwPol"/>
    <cloudIntNetworkP name="IntNetworkP1"/>
  </cloudRouterP>

  <cloudCidr addr="60.10.10.1/16" primary="true">
    <cloudSubnet ip="60.10.10.1/24">
      <cloudRsZoneAttach tDn="uni/clouddomp/provp-aws/region-us-east-1/zone-us-east-1a"/>
    </cloudSubnet>
  </cloudCidr>
</cloudCtxProfile>

<cloudCtxProfile name="prod-payment-east-1" status="">
  <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-aws/region-us-east-1"/>
  <cloudRsToCtx tnFvCtxName="prod-2" status=""/>
  <cloudRouterP name="RouterP1" type="vpn-gw">
    <cloudRsToVpnGwPol tnCloudVpnGwPolName="VgwPol"/>
    <cloudIntNetworkP name="IntNetworkP1" status=""/>
  </cloudRouterP>

  <cloudCidr addr="70.10.10.1/16" primary="true" status="">
    <cloudSubnet ip="70.10.10.1/24" status="">
      <cloudRsZoneAttach tDn="uni/clouddomp/provp-aws/region-us-east-1/zone-us-east-1a"/>
    </cloudSubnet>
  </cloudCidr>
</cloudCtxProfile>

</fvTenant>
</polUni>

```

Step 2 To create a cloud context profile using the **cloud** availability zones supported beginning with release 25.0(2), enter a post such as the following example.

Beginning with release 25.0(2), if you are creating a cloud context profile in a **user** tenant, you are restricted to only **cloud** availability zones. The cloud availability zones are created through the `zone` field highlighted below. For more information on the cloud availability zones, see [Availability Zones, on page 32](#).

Example:

```

<polUni>
<fvTenant name="Corp1" status="">
  <cloudAwsProvider accessKeyId="" secretAccessKey="" providerId="aws" status="" accountId=""/>

  <fvCtx name="prod-1" status="">
    <bgpRtTargetP af="ipv4-ucast">
      <bgpRtTarget rt="route-target:as4-nn2:400:400" type="export"/>
      <bgpRtTarget rt="route-target:as4-nn2:400:400" type="import"/>
    </bgpRtTargetP>
  </fvCtx>
</fvTenant>
</polUni>

```

```

    </bgpRtTargetP>
  </fvCtx>

  <fvCtx name="prod-2" status="">
    <bgpRtTargetP af="ipv4-ucast">
      <bgpRtTarget rt="route-target:as4-nn2:500:500" type="export"/>
      <bgpRtTarget rt="route-target:as4-nn2:500:500" type="import"/>
    </bgpRtTargetP>
  </fvCtx>

  <cloudVpnGwPol name="VgwPol" status=""/>

  <cloudApp name="payment" status="">
    <cloudEPg name="web" status="">
      <cloudRsCloudEPgCtx tnFvCtxName="prod-1" />
    </cloudEPg>
  </cloudApp>
  <cloudApp name="billing">
    <cloudEPg name="app">
      <cloudRsCloudEPgCtx tnFvCtxName="prod-2" />
    </cloudEPg>
  </cloudApp>

  <cloudCtxProfile name="prod-web-east-1">
    <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-aws/region-us-east-1"/>
    <cloudRsToCtx tnFvCtxName="prod-1"/>
    <cloudRouterP name="RouterP1" type="vpn-gw">
      <cloudRsToVpnGwPol tnCloudVpnGwPolName="VgwPol"/>
      <cloudIntNetworkP name="IntNetworkP1"/>
    </cloudRouterP>
    <cloudCidr addr="10.10.0.0/16" primary="yes">
      <cloudSubnet ip="10.10.1.0/24" usage="gateway" scope="public" zone="us-west-1a"/>
      <cloudSubnet ip="10.10.2.0/24" scope="public" zone="us-west-1b"/>
    </cloudCidr>
  </cloudCtxProfile>

  <cloudCtxProfile name="prod-payment-east-1" status="">
    <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-aws/region-us-east-1"/>
    <cloudRsToCtx tnFvCtxName="prod-2" status=""/>
    <cloudRouterP name="RouterP1" type="vpn-gw">
      <cloudRsToVpnGwPol tnCloudVpnGwPolName="VgwPol"/>
      <cloudIntNetworkP name="IntNetworkP1" status=""/>
    </cloudRouterP>
    <cloudCidr addr="20.10.0.0/16" primary="yes">
      <cloudSubnet ip="20.10.1.0/24" scope="public" zone="us-west-1a"/>
    </cloudCidr>
  </cloudCtxProfile>

</fvTenant>
</polUni>

```

Managing a Cloud Region Using the REST API

This section demonstrates how to manage a cloud region using the REST API.

To create a cloud region:

```
<polUni>
```

```

<cloudDomP name="dom-us-east-2">
  <cloudBgpAsP asn="64513"/>
  <cloudProvP vendor="aws">
    <cloudRegion name="us-east-2" adminSt="managed">
      <cloudZone name="us-east-2a"/>
      <cloudZone name="us-east-2b"/>
    </cloudRegion>
  </cloudProvP>
</cloudDomP>
</polUni>

```

Creating a Filter Using the REST API

This section demonstrates how to create a filter using the REST API.

To create a filter:

```

https://<IP_Address>/api/node/mo/.xml
<polUni>
<fvTenant name="intervpc" >
<fvCtx name="VRF1"/>
  <cloudApp name="CloudAP1" >
    <cloudEPg name="CloudEPG1" >
      <cloudRsCloudEPgCtx tnFvCtxName="VRF1"/>
      <fvRsProv tnVzBrCPName="Contract2" > </fvRsProv>
      <cloudEPSelector name="sell" matchExpression="custom:epgtag=='cloudepg1'" />
    </cloudEPg>
  </cloudApp>

  <vzFilter name="http" annotation="orchestrator:msc" >
    <vzEntry name="Entry3" prot="tcp" etherT="ipv4" arpOpc="unspecified" stateful="no"
applyToFrag="no" sFromPort="unspecified" sToPort="unspecified" dFromPort="80" dToPort="80" > </vzEntry>

  </vzFilter>

<vzBrCP name="Contract2" scope="global">
  <vzSubj name="test-subj" >

    <vzRsSubjFiltAtt action="permit" tnVzFilterName="http" directives="none" />

  </vzSubj>
</vzBrCP>
</fvTenant>
</polUni>

```

Creating an Application Profile Using the REST API

This section demonstrates how to create an application profile using the REST API.

Before you begin

Create a tenant.

To create an application profile:

```
https://<IP_Address>/api/node/mo/.xml
<polUni>
<fvTenant name="intervpc" >
<fvCtx name="VRF1"/>
  <cloudApp name="CloudAP1" >

  <cloudEPg name="CloudEPG1" >
    <cloudRsCloudEPgCtx tnFvCtxName="VRF1"/>
    <fvRsProv tnVzBrCPName="Contract2" > </fvRsProv>
    <cloudEPSelector name="sell" matchExpression="custom:epgtag=='cloudepg1' " />
  </cloudEPg>

  </cloudApp>

  <vzFilter name="http" annotation="orchestrator:msc" >
  <vzEntry name="Entry3" prot="tcp" etherT="ipv4" arpOpc="unspecified" stateful="no"
applyToFrag="no" sFromPort="unspecified" sToPort="unspecified" dFromPort="80" dToPort="80" > </vzEntry>

  </vzFilter>
  <vzBrCP name="Contract2" scope="global">
    <vzSubj name="test-subj" >
      <vzRsSubjFiltAtt action="permit" tnVzFilterName="http" directives="none" />
    </vzSubj>
  </vzBrCP>
</fvTenant>
</polUni>
```

Creating a Cloud EPG Using the REST API

This example demonstrates how to create a cloud EPG using the REST API.

Before you begin

Create an application profile and a VRF.

To create a cloud EPG:

Example:

```
<polUni>
<fvTenant name="t2" status="">
  <!-- Tenant provide AWS credentials -->
  <cloudAwsProvider region="us-east-2" accessKeyId="123" secretAccessKey="ABCDE" providerId="admin"/>

  <fvCtx name="v1" status=""/>
  <cloudApp name="ap">
    <cloudEPg name="provEPG" status="">
      <cloudRsCloudEPgCtx tnFvCtxName="v1"/>
      <cloudEPSelector name="1" matchExpression="custom:tag=='provfoo'"/>
    </cloudEPg>
  </cloudApp>
</fvTenant>
</polUni>
```

```

    <cloudEPSelector name="2" matchExpression="custom:tag=='provbaz'"/>
    <fvRsProv tnVzBrCPName="httpFamily"/>
  </cloudEPg>
  <cloudEPg name="consEPG">
    <cloudRsCloudEPgCtx tnFvCtxName="v1"/>
    <cloudEPSelector name="1" matchExpression="custom:tag=='consfoo'"/>
    <cloudEPSelector name="2" matchExpression="custom:tag=='consbaz'"/>
    <fvRsCons tnVzBrCPName="httpFamily"/>
  </cloudEPg>
</cloudApp>
</fvTenant>
</polUni>

```

Creating an External Cloud EPG Using the REST API

This example demonstrates how to create an external cloud EPG using the REST API.

Before you begin

Create an application profile and a VRF.

To create an external cloud EPG:

Example:

```

<polUni>
  <fvTenant name="t2" status="">
    <!-- Tenant provide AWS credentials -->
    <cloudAwsProvider region="us-east-2" accessKeyId="123" secretAccessKey="ABCDE" providerId="admin"/>

    <fvCtx name="v1" status="">
      <cloudApp name="ap">
        <cloudEPg name="provEPGInternet" status="">
          <cloudRsCloudEPgCtx tnFvCtxName="v1"/>
          <cloudEPSelector name="1" matchExpression="custom:tag=='provfoo'"/>
          <cloudEPSelector name="2" matchExpression="custom:tag=='provbaz'"/>
          <fvRsProv tnVzBrCPName="httpFamily"/>
        </cloudEPg>
        <cloudExtEPg name="consInternetEPG">
          <cloudRsCloudEPgCtx tnFvCtxName="v1"/>
          <cloudExtEPSelector name="1" subnet="0.0.0.0/0"/>
          <fvRsCons tnVzBrCPName="httpFamily"/>
        </cloudExtEPg>
      </cloudApp>
    </fvTenant>
  </polUni>

```

Creating a Cloud Template Using the REST API

This section demonstrates how to create a cloud template using the REST API. For more information about cloud templates, see [About the Cloud Template, on page 39](#).

The REST API will change depending on the type of Licensing model selected. The license type of the Cisco Catalyst 8000V is captured by the property `routerThroughput` in the `cloudtemplateProfile` managed object .

If the `routerThroughput` value belongs to **T0/T1/T2/T3** then **BYOL** Cisco Catalyst 8000V is deployed on Cisco Cloud APIC. If `routerThroughput` value is **PAYG** then **PAYG** Cisco Catalyst 8000V is deployed on Cisco Cloud APIC.

Step 1 To create a cloud template post to deploy a **BYOL** Cisco Catalyst 8000V:

```
<polUni>
  <fvTenant name="infra">
    <cloudtemplateInfraNetwork name="default" vrfName="overlay-1">
      <cloudtemplateProfile name="default" routerUsername="admin" routerPassword="rtpssw"
routerThroughput="15"
      routerLicenseToken="hYjZhYjItYTg0mrtrL15ocStS%0AUzRSz0%3"
routerMgmtInterfacePublicIp="yes" routerDataInterfacePublicIp="yes"/>

      <cloudtemplateExtSubnetPool subnetpool="10.20.0.0/16"/>

      <cloudtemplateIntNetwork name="default">
        <cloudRegionName provider="aws" region="us-west-1"/>
        <cloudRegionName provider="aws" region="us-west-2"/>
      </cloudtemplateIntNetwork>

      <cloudtemplateExtNetwork name="default">
        <cloudRegionName provider="aws" region="us-west-2"/>

      <cloudtemplateVpnNetwork name="default">

        <cloudtemplateIpSecTunnel peeraddr="23.2.1.1/32" />
        <cloudtemplateIpSecTunnel peeraddr="23.0.1.1/32" />
        <cloudtemplateIpSecTunnel peeraddr="23.1.1.1/32" />

        <cloudtemplateOspf area="0.0.0.1"/>

      </cloudtemplateVpnNetwork>

      <cloudtemplateBgpEvpn peeraddr="34.1.1.1/32" asn="63000" siteId="123" password="abcd1234" />

    </cloudtemplateExtNetwork>
  </cloudtemplateInfraNetwork>
</fvTenant>
</polUni>
```

Note Beginning with release 25.0(3), tier2 (T2) is the default throughput supported by Cisco Cloud APIC, which is indicated by the property `routerThroughput` in the `cloudtemplateProfile` managed object above.

Step 2 To create a cloud template post to deploy a **PAYG** Cisco Catalyst 8000V:

```
<polUni>
  <fvTenant name="infra">
    <cloudtemplateInfraNetwork name="default" vrfName="overlay-1">
      <cloudtemplateProfile name="default" routerUsername="admin" routerPassword="rtpssw"
routerThroughput="PAYG"
      vmName="c5.4xlarge" routerMgmtInterfacePublicIp="yes" routerDataInterfacePublicIp="yes"/>

      <cloudtemplateExtSubnetPool subnetpool="10.20.0.0/16"/>
```

```

<cloudtemplateIntNetwork name="default">
  <cloudRegionName provider="aws" region="us-west-1"/>
  <cloudRegionName provider="aws" region="us-west-2"/>
</cloudtemplateIntNetwork>

<cloudtemplateExtNetwork name="default">
  <cloudRegionName provider="aws" region="us-west-2"/>

  <cloudtemplateVpnNetwork name="default">

    <cloudtemplateIpSecTunnel peeraddr="23.2.1.1/32" />
    <cloudtemplateIpSecTunnel peeraddr="23.0.1.1/32" />
    <cloudtemplateIpSecTunnel peeraddr="23.1.1.1/32" />

    <cloudtemplateOspf area="0.0.0.1"/>

  </cloudtemplateVpnNetwork>

  <cloudtemplateBgpEvpn peeraddr="34.1.1.1/32" asn="63000" siteId="123" password="abcd1234" />

</cloudtemplateExtNetwork>
</cloudtemplateInfraNetwork>
</fvTenant>
</polUni>

```

On selecting PAYG throughput the user must also select the **vmName** from a list of vmNames which is created by Cloud APIC and represented by the managed object `vmName`.

The following table lists the vmNames that are indicated by the property `vmName` in the `cloudtemplateProfile`.

AWS EC2 Instance	CCR Throughput	vCPUs	Memory
c5.xlarge	up to 5 Gigabit throughput	4	8 GiB
c5.2xlarge	up to 10 Gigabit throughput	8	16 GiB
c5.4xlarge	up to 10 Gigabit throughput	16	32 GiB
c5.9xlarge	up to 10 Gigabit throughput	36	72 GiB
c5n.xlarge	up to 25 Gigabit throughput	4	10.5 GiB
c5n.2xlarge	up to 25Gigabit throughput	8	21 GiB
c5n.4xlarge	up to 25 Gigabit throughput	16	42 GiB
c5n.9xlarge	up to 50 Gigabit throughput	36	96 GiB

Configuring VRF Leak Routes Using the REST API

Before you begin

Review the information provided in [Route Leaking Between Internal VRFs, on page 9](#) and [Global Inter-VRF Route Leak Policy, on page 10](#) before proceeding with the instructions in this section.

Step 1 Enter a post similar to the following to enable or disable contract-based routing.

```
<fvTenant name="infra">
  <cloudVrfRouteLeakPol name="default" allowContractBasedRouting="true"/>
</fvTenant>
```

Where the `allowContractBasedRouting` field has either of the following settings:

- **true:** Indicates that routes are leaked based on contracts in the absence of route maps. When enabled, contracts drive routing when route maps are not configured. When route maps exist, route maps always drives routing.
- **false:** Default setting. Indicates that routes are not leaked based on contracts, and are leaked based on route maps instead.

Step 2 Enter a post similar to the following to use the `leakInternalPrefix` field to configure route leaking for all cloud CIDRs associated with the VRFs.

```
<fvTenant name="t1">
  <fvCtx name="v1">
    <leakRoutes>
      <leakInternalPrefix ip="0.0.0.0/0" le="32">
        <leakTo tenantName="t2" ctxName="v2" scope="public"/>
      </leakInternalPrefix>
    </leakRoutes>
  </fvCtx>
</fvTenant>

<fvTenant name="t2">
  <fvCtx name="v2">
    <leakRoutes>
      <leakInternalPrefix ip="0.0.0.0/0" le="32">
        <leakTo tenantName="t1" ctxName="v1" scope="public"/>
      </leakInternalPrefix>
    </leakRoutes>
  </fvCtx>
</fvTenant>
```

Step 3 Enter a post similar to the following to use the `leakInternalSubnet` field to leak specific routes between a pair of VRFs.

```
<fvTenant name="anyTenant" status="">
  <fvCtx name="VRF1" >
    <leakRoutes status="">
      <leakInternalSubnet ip="110.110.1.0/24" >
        <leakTo ctxName="VRF2" scope="public" tenantName=" anyTenant " />
      </leakInternalSubnet>
    </leakRoutes>
  </fvCtx>
  <fvCtx name="VRF2" status="" >
    <leakRoutes status="">
      <leakInternalSubnet ip="110.110.2.0/24" >
        <leakTo ctxName="VRF1" scope="public" tenantName=" anyTenant " />
      </leakInternalSubnet>
    </leakRoutes>
  </fvCtx>
</fvTenant>
```

Configuring the Source Interface Selection for Tunnels Using the REST API

Before you begin

Review the information provided in [Source Interface Selection for Tunnels, on page 11](#) before proceeding with these instructions.

Enter a post similar to the following to configure the source interface selection for tunnels.

```
<cloudtemplateInfraNetwork name="default" vrfName="overlay-1">
  <cloudtemplateProfile name="defaultxyz" routerUsername="james" routerPassword="bond@@7" />

  <cloudtemplateIpSecTunnelSubnetPool subnetpool="10.20.0.0/16" poolname="pool1" />

  <cloudtemplateIntNetwork name="default">
    <cloudRegionName provider="aws" region="us-west-1"/>
    <cloudRegionName provider="aws" region="us-west-2"/>
  </cloudtemplateIntNetwork>

  <cloudtemplateExtNetwork name="something" vrfName="xyz" >
    <cloudRegionName provider="aws" region="us-west-2"/>
    <cloudtemplateVpnNetwork name="default">
      <cloudtemplateIpSecTunnel peeraddr="23.2.1.1/32" poolname="" presharedkey="abcd"
ikeVersion="v1|v2">
        <b><cloudtemplateIpSecTunnelSourceInterface sourceInterfaceId="2" /></b>
      </cloudtemplateIpSecTunnel>
    </cloudtemplateVpnNetwork>
  </cloudtemplateExtNetwork>
</cloudtemplateInfraNetwork>
```



CHAPTER 5

Viewing System Details

- [Monitoring VM Host Metrics, on page 123](#)
- [Viewing Application Management Details, on page 126](#)
- [Viewing Cloud Resource Details, on page 127](#)
- [Viewing Operations Details, on page 128](#)
- [Viewing Infrastructure Details, on page 130](#)
- [Viewing Administrative Details, on page 130](#)
- [Viewing Health Details Using the Cisco Cloud APIC GUI, on page 132](#)

Monitoring VM Host Metrics

Beginning with release 25.0(1), support is available for monitoring metrics for the VM host where the Cisco Cloud APIC is deployed using the Prometheus Node Exporter. The Prometheus Node Exporter provides visibility to a wide variety of hardware and kernel-related metrics, where it collects technical information from Linux nodes, such as CPU, disk, and memory statistics. For overview information on the Prometheus Node Exporter, see:

<https://prometheus.io/docs/introduction/overview/>

If your Cisco Cloud APIC is running on release 25.0(1) or later, the Prometheus Node Exporter is automatically available by default.


Guidelines and Limitations

HTTP is not supported for monitoring metrics using the Prometheus Node Exporter. Only HTTPS is supported for monitoring metrics using the Prometheus Node Exporter.

Monitoring VM Host Metrics Using the GUI

These procedures describe how to enable the Prometheus Node Exporter to monitor VM host metrics using the GUI.

-
- Step 1** In the Cisco Cloud APIC GUI, navigate to **Infrastructure > System Configuration**, then click on the **Management Access** tab.
- Step 2** In the **HTTPS** area to the right of the window, note the entry in the **Node Exporter** field.




- **Enabled:** The Prometheus Node Exporter has already been enabled. You do not have to continue with these instructions in that case.
- **Disabled:** The Prometheus Node Exporter is not enabled yet. Proceed with these instructions to enable the Prometheus Node Exporter.

Step 3 Click the pencil icon in the **HTTPS** area to edit the HTTPS settings.

The **HTTPS Settings** window appears.

Step 4 Locate the **Node Exporter** field and click **Enable**.




A warning message appears, telling you that saving these settings will restart the web service, and that it will take a moment for it to resume responding to requests. Click **OK** to confirm these changes.

Step 5 At the bottom of the window, click **Save**.

You are returned to the **System Configuration/Management Access** window. The web service reboots and comes back online in a few seconds.

Step 6 In the **HTTPS** area to the right of the window, verify that the entry in the **Node Exporter** field is set to **Enabled**. This verifies that the Prometheus Node Exporter is enabled.



Step 7 Click the link under the **Enabled** text in the **Node Exporter** area.

Another tab in your browser appears, showing the metrics for the VM host where the Cisco Cloud APIC is deployed.

Monitoring VM Host Metrics Using the REST API

These procedures describe how to enable the Prometheus Node Exporter to monitor VM host metrics using the REST API.

Step 1 To determine if the Prometheus Node Exporter is enabled or not, send the following GET call:

```
GET https://<cloud-apic-ip-address>/api/mo/uni/fabric/comm-default/https.xml
```

Locate the `nodeExporter` field to determine if it is set to `enabled` or `disabled`.

Step 2 To monitor VM host metrics, send the following post to enable the Prometheus Node Exporter:

```
POST https://<cloud-apic-ip-address>/api/mo/uni/fabric/comm-default/https.xml
<commHttps nodeExporter="enabled" />
```

The metrics are displayed for the VM host where the Cisco Cloud APIC is deployed.

Step 3 To view the metrics using REST API, send the following GET call:

```
GET https://<cloud-apic-ip-address>/nodeexporter/metrics
```

Step 4 To disable the Prometheus Node Exporter, send the following post:

```
POST https://<cloud-apic-ip-address>/api/mo/uni/fabric/comm-default/https.xml
<commHttps nodeExporter="disabled" />
```

Viewing Application Management Details

This section explains how to view application management details using the Cisco Cloud APIC GUI. The application management details include the information of a specific tenant, application profile, EPG, contract, filter, VRF, service, or cloud context profile.

Step 1 From the **Navigation** menu, choose the **Application Management** tab.

When the **Application Management** tab expands, a list of subtab options appear. See the *Application Management Options* table for more information.

Table 29: Application Management Subtabs

Subtab Name	Description
Tenants	Displays tenants as rows in a summary table.
Application Profiles	Displays application profiles as rows in a summary table.
EPGs	Displays an EPGs as rows in a summary table.
Contracts	Displays a contracts as rows in a summary table.
Filters	Displays filters as rows in a summary table.
VRFs	Displays VRFs as rows in a summary table.
Services	Contains the following two subtabs and information: <ul style="list-style-type: none"> • Devices—Displays the devices as rows in a summary table. • Service Graphs—Displays service graphs as rows in a summary table.
Cloud Context Profiles	Displays cloud context profiles as rows in a summary table.

Step 2 Click the tab that represents the component with the details you want to view.

A summary table appears with items as rows in the table. For example, if you chose the **Tenants** subtab, a list of tenants appear as rows in a summary table

You can filter the rows by clicking the Filter by Attributes bar. Choose the attribute, operator and filter-value. For example, for filtering based on a tenant, choose Tenant == T1 (where T1 is the name of a tenant).

Step 3 To view a summary pane, click the row that represents the specific component you want to view.

Step 4 For more information, double-click the summary table row that represents the specific component you want to view.

A new dialog box appears over the **work** pane with any of the following tabs:

Note The tabs that appear differ between components and configurations.

- **Overview**—Provides a general overview of cloud resources, configuration relationships, and settings of the component.

- **Cloud Resources**—Contains a list of subtabs that display the cloud resource information related to the component.
- **Configuration**—Contains one or more subtabs that display the configuration information related to the component.
- **Statistics**—Enables you to view statistics based on a chosen sampling interval and statistics type. The **Statistics** tab may contain subtabs, depending on the component you are viewing.
- **Event Analytics**—Contains a list of subtabs that display faults, events, and audit logs.

Note The dialog box that appears over the **work** pane contains an **edit** button in the top-right corner between the **refresh** button and the **Actions** button. When clicked, the **edit** button enables you to edit the chosen component.

Viewing Cloud Resource Details

This section explains how to view cloud resource details using the Cisco Cloud APIC GUI. The cloud resource details include the information about a specific region, availability zone, VPC, router, security group, endpoint, instance, and cloud service.

Step 1 From the **Navigation** menu, choose the **Cloud Resources** tab.

When the **Cloud Resources** tab expands, a list of subtab options appear. See the *Cloud Resource Options* table for more information.

Table 30: Cloud Resource Subtabs

Subtab Name	Description
Regions	Displays regions as rows in a summary table.
Availability Zones	Displays the availability zones as rows in a summary table.
VPCs	Displays VPCs as rows in a summary table.
Routers	Displays routers as rows in a summary table.
Security Groups	Displays security groups as rows in a summary table.
Endpoints	Displays endpoints as rows in a summary table.
Instances	Displays the instances as rows in a summary table.
Cloud Services	Contains the following subtabs: <ul style="list-style-type: none"> • Cloud Services Tab—Displays cloud services as rows in a summary table. • Target Groups Tab—Displays target groups as rows in a summary table.

Step 2 Click the tab that represents the component with the details you want to view.

A summary table appears with items as rows in the table. For example, if you chose the **Endpoints** subtab, a list of endpoints appear as rows in a summary table

You can filter the rows by selecting an attribute from the drop-down menu when you click the *Filter by attributes* bar. The attributes displayed in the drop-down menu depend on the selected subtab.

Step 3 To view a summary pane, click the row that represents the specific component you want to view.

Step 4 For more information, double-click the summary table row that represents the specific component you want to view.

A new dialog box appears over the **work** pane with any of the following tabs:

Note The tabs that appear differ between components and configurations.

- **Overview**—Provides a general overview of cloud resources, configuration relationships, and settings of the component.
- **Cloud Resources**—Contains a list of subtabs that display the cloud resource information related to the component.
- **Application Management**—Contains a list of subtabs that display the ACI relation information related to the component.
- **Statistics**—Enables you to view statistics based on a chosen sampling interval and statistics type. The **Statistics** tab may contain subtabs, depending on the component you are viewing.
- **Event Analytics**—Contains a list of subtabs that display faults, events, and audit logs.

Viewing Operations Details

This section explains how to view operations details using the Cisco Cloud APIC GUI. The operations details include the information of a specific fault, event, audit log, active sessions, backup and restore policies, tech support policies, firmware management, scheduler policies, and remote locations.

Step 1 From the **Navigation** menu, choose the **Operations** tab.

When the **Operations** tab expands, a list of subtab options appear. See the *Operations Options* table for more information.

Table 31: Operations Subtabs

Subtab Name	Description
Event Analytics	Contains the following subtabs: <ul style="list-style-type: none"> • Faults Tab—Displays faults as rows in a summary table. • Events Tab—Displays events as rows in a summary table. • Audit Logs Tab—Displays audit logs as rows in a summary table.
Active Sessions	Displays a list of active users as rows in a summary table.

Subtab Name	Description
Backup & Restore	Contains the following subtabs: <ul style="list-style-type: none"> • Backups Tab—Displays backups as rows in a summary table. • Backup Policies Tab—Displays backup policies as rows in a summary table. • Job Status Tab—Displays the job status as rows in a summary table. • Event Analytics Tab—Contains the following subtabs: <ul style="list-style-type: none"> • Faults Tab—Displays faults as rows in a summary table. • Events Tab—Displays events as rows in a summary table. • Audit Logs Tab—Displays audit logs as rows in a summary table.
Tech Support	Contains the following subtabs: <ul style="list-style-type: none"> • Tech Support Tab—Displays tech support policies as rows in a summary table. • Core Logs Tab—Displays core logs as rows in a summary table. • Per-Feature Containers Tab—Displays the per-feature containers as rows in a summary table.
Firmware Management	Contains the following subtabs: <ul style="list-style-type: none"> • General Tab—Displays general firmware management information. • Images Tab—Displays a list of images. • Event Analytics Tab—Contains the following subtabs: <ul style="list-style-type: none"> • Faults Tab—Displays faults as rows in a summary table. • Events Tab—Displays events as rows in a summary table. • Audit Logs Tab—Displays audit logs as rows in a summary table.
Schedulers	Displays scheduler policies as rows in a summary table.
Remote Locations	Displays remote locations as rows in a summary table.

Step 2 Click the tab that represents the component you want to view.

A summary table appears with items as rows in the table. For example, if you chose the **Active Sessions** subtab, a list of active sessions appear as rows in a summary table.

You can filter the rows by clicking the *Filter by Attributes* bar. Choose the attribute, operator and filter-value. For example, for filtering based on a username, choose username == user1 (where user1 is a user logged into Cloud APIC).

Step 3 To view a summary pane, click the row that represents the specific component you want to view.

Step 4 For more information, double-click the summary table row that represents the specific item you want to view.

A new dialog box appears over the **work** pane that displays additional information about the item you chose from the summary table.

Viewing Infrastructure Details

This section explains how to view infrastructure details using the Cisco Cloud APIC GUI. The infrastructure details include information about system configuration, inter-region connectivity, and external connectivity.

Step 1 From the **Navigation** menu, choose the **Infrastructure** tab.

When the **Infrastructure** tab expands, a list of subtab options appear. See the *Infrastructure Options* table for more information.

Table 32: Infrastructure Subtabs

Subtab Name	Description
System Configuration	Displays General system configuration information, Management Access information, Controllers , and Event Analytics .
Inter-Region Connectivity	Displays one pane with a map that contains the inter-region connectivity view and additional panes for each region.
Inter-Site Connectivity	Displays one pane with a map that contains the inter-site connectivity view and additional panes for each region.

Step 2 Click the tab that represents the component with the details you want to view.

Viewing Administrative Details

This section explains how to view administrative details using the Cisco Cloud APIC GUI. The administrative details include the information about authentication, security, users, and smart licensing..

Step 1 From the **Navigation** menu, choose the **Administrative** tab.

When the **Administrative** tab expands, a list of subtab options appear. See the *Administrative Options* table for more information.

Table 33: Administrative Subtabs

Subtab Name	Description
Authentication	<p>Displays the Authentication Default Settings, Login Domains, and Providers subtabs, which contain the information described below:</p> <ul style="list-style-type: none"> • Authentication Default Settings Tab—Displays settings information. • Login Domains Tab—Displays the login domains as rows in a summary table. • Providers Tab—Displays the providers as rows in a summary table. • Event Analytics Tab—Displays the Faults, Events, and Audit Logs subtabs, each with the corresponding information displayed as rows in a summary table.
Security	<p>Contains the following list of subtabs:</p> <ul style="list-style-type: none"> • Security Default Settings Tab—Enables you to view the default security settings information. • Security Domains Tab—Enables you to view security domain information in a summary table. • Roles Tab—Enables you to view the role information in a summary table. • RBAC Rules Tab—Enables you to view RBAC rule information in a summary table. • Certificate Authorities Tab—Enables you to view the certificate authority information in a summary table. • Key Ring Tab—Enables you to view key ring information in a summary table.
Users	<p>Contains the following subtabs:</p> <ul style="list-style-type: none"> • Local Tab—Displays local users as rows in a summary table. • Remote Tab—Displays remote users as rows in a summary table.

Subtab Name	Description
Smart Licensing	<p>Contains the following subtabs:</p> <ul style="list-style-type: none"> • General Tab—Displays the licenses as rows in a summary table. • Faults Tab—Displays faults as rows in a summary table.

Step 2 Click the tab that represents the component you want to view.

For some options, a summary table appears with items as rows in the table (For example, if you choose the **Users** tab, a list of users appear as rows in a summary table). To view a summary pane, click the row that represents the specific component you want to view. To view more information, double-click the summary table row that represents the specific item you want to view. A new dialog box appears over the **work** pane that displays additional information about the item you chose from the summary table.

Note You can filter the rows by entering an attribute in the *Filter by Attributes* bar.

Viewing Health Details Using the Cisco Cloud APIC GUI

This section explains how to view health details using the Cisco Cloud APIC GUI. You can view health details for any object that you can see in the Cloud Resources area in the Cisco Cloud APIC GUI, such as the following:

- Regions
- Availability Zones (for AWS cloud sites)
- VPCs (for AWS cloud sites)
- VNETs (for Azure cloud sites)
- Routers
- Security Groups
- Endpoints
- Instances
- Cloud Services

Step 1 From the **Navigation** menu, choose the **Dashboard** tab.

The **Dashboard** window for the Cisco Cloud APIC system appears. From this window, you can view the overall health status of your system.

The screenshot shows the Cisco Cloud APIC GUI Dashboard. The left sidebar contains navigation menus for Dashboard, Application Management, Cloud Resources, Operations, Infrastructure, and Administrative. The main content area is titled "Dashboard" and displays several key metrics:

- Health Summary:** A large orange box labeled "Major" indicates the overall system health.
- Fault Summary:** A bar chart showing fault counts by severity: Critical (2), Major (14), Minor (4), and Warning (2).
- Inter-Site Connectivity Status:** Shows 4 CSRs, 4 IPsec Tunnels, 4 OSPF, and 0 BGP Sessions.
- Inter-Region Connectivity Status:** Shows 4 CSRs, 0 Virtual Networks, 0 IPsec Tunnels, and 0 BGP Sessions.
- Smart License Registration State:** Shows "Unregistered".
- Smart License Authorization Status:** Shows "Evaluation" with 76 days remaining.
- Cloud Resources Summary (Azure):**
 - Regions: 0 Total
 - Virtual Networks: 2 Total
 - Routers: 4 Total
 - Endpoints: 0
 - Virtual Machines: 0

Step 2 Click within the Fault Summary area in the **Dashboard** window.

The **Event Analytics** window appears, showing more detailed information for the specific fault level that you clicked. The following screen shows an example **Event Analytics** window for the faults listed with critical severity.

The screenshot shows the Cisco Cloud APIC GUI Event Analytics window. The left sidebar is the same as in the previous screenshot. The main content area is titled "Event Analytics" and displays a table of faults filtered by "Severity: Critical".

Acked	Severity	Code	Affected object	Description	Lifecycle	Creation Time
<input type="checkbox"/>	Critical	F0104	topology/pod-1/node-1/sys/caggr-[po1.1]	Bond Interface po1.1 on node 1 of fabric mininet with hostname capic1 is now down	raised	Sep 11 2019 05:22:33pm
<input type="checkbox"/>	Critical	F0104	topology/pod-1/node-1/sys/caggr-[po1]	Bond Interface po1 on node 1 of fabric mininet with hostname capic1 is now down	raised	Sep 11 2019 05:22:33pm

At the bottom of the table, there is a pagination control showing "10 Rows" and "Page 1 of 1".

Step 3 Click the **X** next to the Severity level to display Event Analytics information for all faults.

The information provided in the **Event Analytics** window changes to show the events with critical, major, and warning levels of severity.

Viewing Health Details Using the Cisco Cloud APIC GUI

Acked	Severity	Code	Affected object	Description	Lifecycle	Creation Time
No	Critical	F0104	topology/pod-1/node-1/ryscapgr-[po1.1]	Bond Interface po1.1 on node 1 of fabric mininet with hostname capic1 is now down	raised	Sep 11 2019 05:22:33pm
No	Critical	F0104	topology/pod-1/node-1/ryscapgr-[po1]	Bond Interface po1 on node 1 of fabric mininet with hostname capic1 is now down	raised	Sep 11 2019 05:22:33pm
No	Major	F3442	acct-{infra}/region-{eastus}/context-{overlay-1}-addr-[10.10.0.128/20]/csr-[cl_routerp_eastus_1_0]/instoper	Operational State of the hcloudInstanceOper is down with [compute.VirtualMachineClientCreateOrUpdate: Failure sending request: StatusCode=404 -- Original Error: Code= "ResourceGroupNotFound" Message="Resource group 'APFC-infra-mininet-fchazet-centralus' could not be found."]	raised	Sep 11 2019 07:38:27pm
No	Major	F3442	acct-{infra}/region-{centralus}/context-{overlay-1}-addr-[10.10.0.0/25]/csr-[cl_routerp_centralus_1_0]/instoper	Operational State of the hcloudInstanceOper is down with [compute.VirtualMachineClientCreateOrUpdate: Failure sending request: StatusCode=404 -- Original Error: Code= "ResourceGroupNotFound" Message="Resource group 'APFC-infra-mininet-fchazet-centralus' could not be found."]	raised	Sep 11 2019 07:38:27pm
No	Major	F3442	acct-{infra}/region-{eastus}/context-{overlay-1}-addr-[10.10.0.0/25]/csr-[cl_routerp_eastus_0_0]/instoper	Operational State of the hcloudInstanceOper is down with [compute.VirtualMachineClientCreateOrUpdate: Failure sending request: StatusCode=404 -- Original Error: Code= "ResourceGroupNotFound" Message="Resource group 'APFC-infra-mininet-fchazet-centralus' could not be found."]	raised	Sep 11 2019 07:39:27pm
No	Major	F3442	acct-{infra}/region-{centralus}/context-{overlay-1}-addr-[10.10.0.0/25]/csr-[cl_routerp_centralus_0_0]/instoper	Operational State of the hcloudInstanceOper is down with [compute.VirtualMachineClientCreateOrUpdate: Failure sending request: StatusCode=404 -- Original Error: Code= "ResourceGroupNotFound" Message="Resource group 'APFC-infra-mininet-fchazet-centralus' could not be found."]	raised	Sep 11 2019 07:45:10pm
No	Major	F3527	acct-{infra}/region-{eastus}/context-{overlay-1}-addr-[10.10.0.128/20]/csr-[cl_routerp_eastus_0_0]/license/oper	Operational State of the HcpPlatformLicense is down with administrative-down	raised	Sep 11 2019 05:21:24pm
No	Major	F3527	acct-{infra}/region-{centralus}/context-{overlay-1}-addr-[10.10.0.0/25]/csr-[cl_routerp_centralus_1_0]/license/oper	Operational State of the HcpPlatformLicense is down with administrative-down	raised	Sep 11 2019 05:21:35pm
No	Major	F0101	topology/pod-1/node-1/ryscapgr-[dev/vcb]-[1]	Storage unit dev/vcb on node 1 with hostname capic1 has failed.	raised	Sep 11 2019 05:22:33pm

Step 4 From the **Navigation** menu, choose the **Cloud Resources** tab.

When the **Cloud Resources** tab expands, a list of subtab options appear. See the *Administrative Options* table for more information.

Step 5 Choose any item under the **Cloud Resources** tab to display health information for that component.

For example, the following figure shows health information that might be displayed when you click on **Cloud Resources > Regions**, then you select a specific region.

Name	Admin State	Tenants	EPGs	AZs	Virtual Networks
eastus	managed	N/A	N/A	N/A	N/A
eastus2	managed	N/A	N/A	N/A	N/A
westus	managed	N/A	N/A	N/A	N/A
centralus	managed	N/A	N/A	N/A	N/A
koreasouth	unmanaged	N/A	N/A	N/A	N/A
francecentral	unmanaged	N/A	N/A	N/A	N/A
eastasia	unmanaged	N/A	N/A	N/A	N/A
canadeast	unmanaged	N/A	N/A	N/A	N/A
brazilouth	unmanaged	N/A	N/A	N/A	N/A
australiaseast	unmanaged	N/A	N/A	N/A	N/A
australacentral2	unmanaged	N/A	N/A	N/A	N/A
koreacentral	unmanaged	N/A	N/A	N/A	N/A
ukwest	unmanaged	N/A	N/A	N/A	N/A
southindia	unmanaged	N/A	N/A	N/A	N/A
southeastasia	unmanaged	N/A	N/A	N/A	N/A

Cloud Provider's Region eastus

CRITICAL	MAJOR	MINOR	WARNING
0	0	0	0

General

Region: region-eastus

Usage

0 Total

Settings

Admin State: Managed

Oper State: In use

Account: infra

Cloud Provider ID: ct_ctprofile_eastus



CHAPTER 6

Deploying Layer 4 to Layer 7 Services

- [Overview, on page 135](#)
- [Deploying a Service Graph, on page 139](#)

Overview

The Cisco Cloud APIC enables you to deploy Layer 4 to Layer 7 service devices to the public cloud. This initial release supports application load balancer (ALB) deployments in Amazon Web Services (AWS).

About Application Load Balancers

An application load balancer (ALB) is a Layer 7 load balancer that inspects packets and creates access points to HTTP and HTTPS headers. It also identifies the load and spreads it out to the targets with higher efficiency. You deploy an ALB using a service graph, which enables you to define how you want traffic to come into the network, the devices that the traffic passes through, and how the traffic leaves the network. You specify these actions by configuring one or more listeners.

Listeners enable you to specify the ports and protocols (HTTP or HTTPS) that the ALB accepts traffic on. When specifying HTTPS, you also choose a security policy and an SSL certificate.




Note A listener can have multiple certificates.

All listeners require you to configure at least one rule (a default rule, which does not have a condition). Rules enable you to specify the action that the load balancer takes when a condition is met. For example, you can create a rule that redirects traffic to a specified URL when a request is made to a specified hostname or path.


There are two deployment types: internet-facing and internal-facing. An internet-facing deployment inserts the ALB as a service between the consumer external EPG and the provider cloud EPG. The following figure shows the contract configuration within the VRF and the ALB as a service inserted between the consumer external EPG and the provider cloud EPG.

Figure 21: Internet-Facing Deployment



An internal-facing deployment inserts the ALB as a service between the consumer cloud EPG and the provider cloud EPG. The following figure shows the contract configuration within the VRF and the ALB as a service inserted between the consumer cloud EPG and provider cloud EPG.

Figure 22: Internal-Facing Deployment



Note You can find more information about ALBs in the documentation on the AWS website.

Dynamic Server Attachment to Server Pool

Servers in the server pool or target group are dynamically added. You do not need to specify the IP addresses or instance Ids for the targets. The relation from a listener rule to a provider cloud EPG is used for the dynamic selection of endpoints. The relation is also used for adding the endpoints to the target group. By default, the endpoints are registered with the port number 80.

Based on the target group-to-security group association that is provided in the ALB, and the EPG (security group) of the endpoint, the EC2 instance (server) is associated to the target group dynamically on the target group's default port. Alternatively, instead of registering the EC2 instance on the target group port, you can attach the custom port by specifying the ports in the following table:

Table 34: Custom Port-Based Attachment

Provider EPG	Ports
EPGMap:<Epg1DN>	9090

Provider EPG	Ports
EPGMap:<Epg2DN>	9091, 9099

You can specify EPGMap:<EpgDN> as the tag and the list of ports to be registered on the target group as a list separated by commas.

About Service Graphs

The Cisco Application Centric Infrastructure (ACI) treats services as a part of an application. Any services that are required are treated as a service graph that is instantiated on the Cisco ACI fabric from the Cisco APIC. You define the service for the application while service graphs identify the set of network or service functions that the application needs.

A service graph represents the network using the following elements:

- **Function node**—A function node represents a function that is applied to the traffic, such as a load balancer. A function within the service graph might require one or more parameters and have one or more connectors.
- **Terminal node**—A terminal node enables input and output from the service graph.
- **Connector**—A connector enables input and output from a node.

After the graph is configured, the Cisco APIC automatically configures the services according to the service function requirements that are specified in the service graph. The Cisco APIC also automatically configures the network according to the needs of the service function that is specified in the service graph, which does not require any change in the service device.

A service graph is represented as two or more tiers of an application with the appropriate service function inserted between them.

A service appliance (device) performs a service function within the graph. One or more service appliances might be required to render the services required by a graph. A single-service device can perform one or more service functions.

Service graphs and service functions have the following characteristics:

- Traffic sent from specific endpoint groups can be redirected based on a policy.
- Service graph redirection is directional. In other words, redirection can be applied to both traffic directions or either one of them.
- Logical functions can be rendered on the appropriate device, based on the policy.
- The service graph supports splits and joins of edges, and it does not restrict the administrator to linear service chains.
- Traffic can be reclassified again in the network after a service appliance emits it.

By using a service graph, you can install a service, a load balancer, once and deploy it multiple times in different logical topologies. Each time the graph is deployed, Cisco ACI takes care of changing the configuration on the service device to enable the forwarding in the new logical topology.

About Function Nodes

A function node represents a single service function. A function node has function node connectors, which represent the network requirement of a service function.

A function node within a service graph requires the following parameters:

- A tenant
- A cloud context profile with subnets in two availability zones

Function parameters can be specified when the service graph is rendered. For example, if the function node is a load balancer, the listener and its rule can be specified for the function node at the time the graph is rendered.

About Terminal Nodes

Terminal nodes connect a service graph with the contracts. You can insert a service graph for the traffic between two application cloud EPGs by connecting the terminal node to a contract. Once connected, traffic between the consumer cloud EPG and provider cloud EPG of the contract is redirected to the service graph.

Deploying a Service Graph

The service graph enables you to define how traffic flows between devices, how the traffic comes into the network, which devices the traffic passes through, and how the traffic leaves the network.

Before you can configure a service graph, you must first configure the following:

1. A tenant
2. A cloud context profile
3. Subnets
4. An application profile
5. A consumer EPG
6. A provider EPG
7. A contract

Deploying the Service Graph Using the Cloud APIC GUI

Creating a Load Balancer Using the Cisco Cloud APIC GUI

This section explains how to create a load balancer using the Cisco Cloud APIC GUI.

-
- Step 1** Click **Application Management > Services**.
The **Services** page appears.

Step 2 Click the Devices tab, then click **Actions > Create Device**.

The **Create Device** page appears.

Step 3 Enter the appropriate values in each field as listed in the following *Create Device Dialog Box Fields* table then continue.

Table 35: Create Device Dialog Box Fields

Properties	Description
General	
Name	Enter the name of the load balancer.
Tenant	To choose a tenant: <ol style="list-style-type: none"> Click Select Tenant. The Select Tenant dialog appears. From the column on the left, click to choose a tenant. Click Select. You return to the Create Device dialog box.
Settings	
Service Type	Choose Application Load Balancer .
Scheme	Choose Internal or Internet Facing .
Subnets	You can specify only one subnet per availability zone. You must specify subnets from at least two availability zones to increase the availability of your load balancer. <ol style="list-style-type: none"> Click Add Subnet. The Add Subnet dialog box appears. In the Add Subnet dialog box, click Select Cloud Context Profile. The Select Cloud Context Profile dialog box appears. In the Select Cloud Context Profile dialog box, select a cloud context profile, then click Select. You are returned to the Add Subnet dialog box. In the Add Subnet dialog box, click Select Subnet. The Select Subnet dialog box appears. In the Select Subnet dialog box, select a subnet, then click Select. You are returned to the Add Subnet dialog box. In the Add Subnet dialog box, click Add. You are returned to the Create Device page.

Step 4 Click **Save** when finished.

Creating a Service Graph Template Using the Cisco Cloud APIC GUI

This section explains how to configure a service graph template using the Cisco Cloud APIC GUI.

Before you begin

You have already created a device.

Step 1 Click **Application Management > Services**.

The **Services** page appears.

Step 2 Click the **Service Graphs** tab, then click **Actions > Create Service Graph**.

The **Create Service Graph** page appears.

Step 3 Enter the appropriate values in each field as listed in the following *Create Service Graph Dialog Box Fields* table then continue.

Table 36: Create Service Graph Dialog Box Fields

Properties	Description
General	
Name	Enter the name of service graph template.
Tenant	To choose a tenant: <ol style="list-style-type: none"> a. Click Select Tenant. The Select Tenant dialog appears. b. From the column on the left, click to choose a tenant. c. Click Select. You return to the Create Service Graph dialog box.
Description	Enter a description of the service graph template.
Settings	

Properties	Description
Select a Device	<p>To choose a device:</p> <ol style="list-style-type: none"> Drag and drop the Application Load Balancer icon to the Drop Device area in the service graph. The Service Node dialog box appears. Click Select Application Load Balancer. The Select Application Load Balancer dialog appears. From the column on the left, click to choose a device. Click Select. You return to the Service Node dialog box. Click Add. You return to the Create Service Graph window.

Step 4 Click **Save** when finished.

Deploying Layer 4 to Layer 7 Services Using the Cisco Cloud APIC GUI

This section explains how to deploy Layer 4 to Layer 7 services.

Before you begin

- You have configured a device.
- You have configured a service graph.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Configuration**.

A list of **Configuration** options appear in the **Intent** menu.

Step 3 From the **Configuration** list in the **Intent** menu, click **EPG Communication**. The **EPG Communication** dialog box appears with the **Consumer EPGs**, **Contract**, and **Provider EPGs** information.

Step 4 To choose a contract:

- Click **Select Contract**. The **Select Contract** dialog appears.
- In the pane on the left side of the **Select Contract** dialog, click to choose a contract then click **Select**. The **Select Contract** dialog box closes.

Step 5 To add a consumer EPG:

- Click **Add Consumer EPGs**. The **Select Consumer EPGs** dialog appears.
- In the pane on the left side of the **Select Consumer EPGs** dialog, click to place a check in a check box to choose a cloud EPG (for an internal facing load balancer) or a cloud external EPG (for an internet facing load balancer) then click **Select**. The **Select Consumer EPGs** dialog box closes.

- Step 6** To add a provider EPG:
- Click **Add Provider EPGs**. The **Select Provider EPGs** dialog appears.
 - In the pane on the left side of the **Select Provider EPGs** dialog, click to place a check in a check box to choose a provider EPG then click **Select**. The **Select Provider EPGs** dialog box closes.

- Step 7** To choose a service graph:
- From the **EPG Communication Configuration** dialog, click **Select Service Graph**. The **Select Service Graph** dialog box appears.
 - In the pane on the left side of the **Select Service Graph** dialog, click to choose a service graph then click **Select**. The **Select Service Graph** dialog box closes.

- Step 8** Under **Service Graph Preview**, click **Add Cloud Load Balancer Listener**. The **Add Cloud Load Balancer Listener** dialog appears that enables you to add listeners.

Listeners are the ports and protocols that the device will work on.

- Step 9** Enter the appropriate values in each field as listed in the following *Add Cloud Load Balancer Listener Dialog Box Fields* table then continue.

Table 37: Add Cloud Load Balancer Listener Dialog Box Fields

Properties	Description
Name	Enter the name of the listener.
Port	Enter the port that the device will accept traffic on.
Protocol	Click to choose HTTP or HTTPS .
Security Policy	Click the drop-down list and choose a security policy (only available when HTTPS is chosen).
SSL Certificate	<p>To choose an SSL certificate(only available when HTTPS is chosen):</p> <ol style="list-style-type: none"> Click Add SSL Certificates. Click to place a check mark in the check box of the certificates you want to add. Choose a key ring: <ol style="list-style-type: none"> Click Select Key Ring. The Select Key Ring dialog appears. From the Select Key Ring dialog, click to choose a key ring in the left column then click Select. The Select Key Ring dialog box closes. Click the Certificate Store drop-down list and choose a certificate. <p>Note A listener can have multiple certificates.</p>

Properties	Description
Add Rule	To add rule settings to the device listener, click Add Rule . A new row appears in the Rules list an the Rules Settings fields are enabled.

Properties	Description
Rule Settings	

Properties	Description
	<p>The Rule Settings pane contains the following options:</p> <ul style="list-style-type: none"> • Name—Enter a name for the rule. • Host—Enter a hostname to create a host-based condition. When a request is made for this hostname, the action you specify is taken. • Path—Enter a path to create a path-based condition. When a request is made for this path, the action you specify is taken. • Type—The action type tells the device which action to take. The action type options: <ul style="list-style-type: none"> • Return fixed response—Returns a response using the following options: <ul style="list-style-type: none"> • Fixed Response Body—Enter a response message. • Fixed Response Code—Enter a response code. • Fixed response Content-Type—Choose a content type. • Forward—Forwards traffic using the following options: <ul style="list-style-type: none"> • Port—Enter the port that the device will accept traffic on. • Protocol—Click to choose HTTP or HTTPS. • Provider EPG—The EPG with the web server that handles the traffic. • EPG—To choose an EPG: <ol style="list-style-type: none"> a. Click Select EPG. The Select EPG dialog box appears. b. From the Select EPG dialog ox, click to choose an EPG in the left column then click Select. The Select EPG dialog box closes. • Redirect—Redirects requests to another location using the following options: <ul style="list-style-type: none"> • Redirect Code—Click the Redirect Code drop-down list and choose a code.

Properties	Description
	<ul style="list-style-type: none"> • Redirect Hostname—Enter a hostname for the redirect. • Redirect Path—Enter a redirect path. • Redirect Port—Enter the port that the device will accept traffic on. • Redirect Protocol—Click to the Redirect Protocol drop-down list and choose HTTP, HTTPS, or Inherit. • Redirect Query—Enter a redirect query. <p>Click Add Rule when finished.</p>

Step 10 Click **Add** when finished.
The service graph is deployed.

Deploying a Service Graph Using the REST API

Creating an Internal-Facing Load Balancer Using the REST API

This example demonstrates how to create an internal-facing load balancer using the REST API.

To create an internal-facing load balancer:

Example:

```
<polUni>
  <fvTenant name="t2" status="">
    <cloudLB name="ALB1" type="application" scheme="internal" status="">
      <cloudRsLDevToCloudSubnet tDn="uni/tn-t2/ctxprofile-cl/cidr-[10.33.0.0/16]/subnet-[10.33.7.0/24]"
status=""/>
      <cloudRsLDevToCloudSubnet tDn="uni/tn-t2/ctxprofile-cl/cidr-[10.33.0.0/16]/subnet-[10.33.8.0/24]"
status=""/>
    </cloudLB>
  </fvTenant>
</polUni>
```

Configuring an Internet-Facing Load Balancer Using the REST API

This example demonstrates how to create an internet-facing load balancer using the REST API.

To create an internet-facing load balancer:

Example:

```
<polUni>
  <fvTenant name="t2" status="">
    <cloudLB name="ALB1" type="application" scheme="internet" status="">
      <cloudRsLDevToCloudSubnet tDn="uni/tn-t2/ctxprofile-c1/cidr-[10.33.0.0/16]/subnet-[10.33.5.0/24]"
        status=""/>
      <cloudRsLDevToCloudSubnet tDn="uni/tn-t2/ctxprofile-c1/cidr-[10.33.0.0/16]/subnet-[10.33.6.0/24]"
        status=""/>
    </cloudLB>
  </fvTenant>
</polUni>
```

Creating a Service Graph Using the REST API

This example demonstrates how to create a service graph using the REST API.

To create a service graph:

```
<polUni>
  <fvTenant name="t2">
    <vnsAbsGraph name="CloudGraph" type="cloud" status="">
      <vnsAbsTermNodeProv name="Input1">
        <vnsAbsTermConn name="C1"/>
      </vnsAbsTermNodeProv>
      <vnsAbsTermNodeCon name="Output1">
        <vnsAbsTermConn name="C2"/>
      </vnsAbsTermNodeCon>
      <vnsAbsNode funcType="GoTo" name="N1" managed="yes">
        <vnsRsNodeToCloudLDev tDn="uni/tn-t2/clb-ALB1" status=""/>
        <vnsAbsFuncConn name="provider"/>
        <vnsAbsFuncConn name="consumer"/>
      </vnsAbsNode>
      <vnsAbsConnection connDir="consumer" connType="external" name="CON2">
        <vnsRsAbsConnectionConns tDn="uni/tn-t2/AbsGraph-CloudGraph/AbsTermNodeCon-Output1/AbsTConn"/>
        <vnsRsAbsConnectionConns tDn="uni/tn-t2/AbsGraph-CloudGraph/AbsNode-N1/AbsFConn-consumer"/>
      </vnsAbsConnection>
      <vnsAbsConnection connDir="provider" connType="internal" name="CON1">
        <vnsRsAbsConnectionConns tDn="uni/tn-t2/AbsGraph-CloudGraph/AbsTermNodeProv-Input1/AbsTConn"/>
        <vnsRsAbsConnectionConns tDn="uni/tn-t2/AbsGraph-CloudGraph/AbsNode-N1/AbsFConn-provider"/>
      </vnsAbsConnection>
    </vnsAbsGraph>
  </fvTenant>
</polUni>
```

Attaching a Service Graph Using the REST API

This example demonstrates how to attach a service graph using the REST API.

To attach a service graph:

```

<polUni>
  <fvTenant name="t2">
    <vzBrCP name="httpFamily">
      <vzSubj name="default" revFltPorts="yes" targetDscp="unspecified">
        <vzRsSubjGraphAtt tnVnsAbsGraphName="CloudGraph"/>
      </vzSubj>
    </vzBrCP>
  </fvTenant>
</polUni>

```

Configuring an HTTP Service Policy Using the REST API

This example demonstrates how to create an HTTP service policy using the REST API.

To create an HTTP service policy:

```

<polUni>
  <fvTenant name="t2">
    <vnsAbsGraph name="CloudGraph" type="cloud" status="">
      <vnsAbsNode funcType="GoTo" name="N1" managed="yes">
        <cloudSvcPolicy tenantName="t2" contractName="httpFamily" subjectName="consubj">
          <cloudListener name="http_listener1" port="80" protocol="http" status="">
            <cloudListenerRule name="rule1" priority="10" default="yes" status="">
              <cloudRuleAction type="forward" port="80" protocol="http"
epgdn="uni/tn-t2/cloudapp-ap/cloudepg-provEPG"/>
            </cloudListenerRule>
            <cloudListenerRule name="redirectRule" priority="20">
              <cloudRuleCondition type="path" value="/img/*"/>
              <cloudRuleAction type="redirect" RedirectPort="8080"/>
            </cloudListenerRule>
            <cloudListenerRule name="FixedRspRule" priority="30">
              <cloudRuleCondition type="host" value="example.com"/>
              <cloudRuleAction type="fixedResponse" FixedResponseCode="200"/>
            </cloudListenerRule>
            <cloudListenerRule name="redirectHPRule" priority="40" status="">
              <cloudRuleCondition type="host" value="example.com"/>
              <cloudRuleCondition type="path" value="/img/*"/>
              <cloudRuleAction type="forward" port="80" protocol="http"
epgdn="uni/tn-t2/cloudapp-ap/cloudepg-provEPG"/>
            </cloudListenerRule>
          </cloudListener>
        </cloudSvcPolicy>
      </vnsAbsNode>
    </vnsAbsGraph>
  </fvTenant>
</polUni>

```

Configuring a Key Ring Using the REST API

This example demonstrates how to configure a key ring using the REST API. For more information about key ring configuration, see the *Cisco APIC Basic Configuration Guide*.


```

MjIwNTMwNVoXDTE5MTAwMjIwNTMwNVowgY0xCzAJBgNVBAYTA1VTMQswCQYDVQQLI
EwJJDQTERMA8GA1UEBxMIU2FuIEpvc2UxEjAQBGNVBAOTCU15Q29tcGFueTEOMAwG
A1UECmFTX1PcmcxGDAWBGNVBAUDyouYw1hem9uYXdzLmNvbTEgMB4GCSqGSIb3
DQEJARYRcmFtc2hhaEBjaXNjby5jb20wggeiMA0GCSqGSIb3DQEBAQUAA4IBDwAw
ggEKAoIBAQQdMgFor5Ee/+dOgcueYMGryF8uKaBf/M0lAL1sa7OvwyPt2bRe4d9B
ga/SHU+0vg93F/mqMHQ1seMBUHUBDxwOISSABfL0qbbvJKjZ+gqvI2oJF4aKef8
KAXv1A8h53nrX5Jw0Nk+394x4cC5Ff8/KQpRq1ZadwZqeO8epz5I4s8XpMOBDMfA
4ccW/IzYNjxt9lhataYaw7smpnNs8ym0DZBZuUguxKgit2QgiB/pl9jL8Q1sf6dg
3aslPyXNizHPRIzHSHFadOI3Y2INj9lXrfLEJd8uD2qk1kK4Pwo590Jk8Sry1qSJ
YHGJHn8de+xxYBlZCyIqAbWTg0RsUD1AgMBAAGjgUwgfIwHQYDVR0OBBYEFBYq
K3b39+1oOr4IBSsePwcOpML7MIHCBGNVHSMegbowgbeAFBYqK3b39+1oOr4IBSse
PwcOpML7oYGTpIGQMIGNMQswCQYDVQQGEwJVUzELMAkGA1UECBMCQ0EExETAPBgNV
BACTFNhbIBKb3NlMRIwEAYDVQQKEw1NeUNvbXBhbnkxKjJAMBGNVBAStBU15T3Jn
MRgwFgYDVQQDFAF8qLmFtYXNjaXNjby5jb20xIDAeBgkqhkiG9w0BCQEWEXJhbXNo
YWhAY2lzy28uY29tggkApY20n/9qsGwwDAYDVR0TBAAUwAwEB/zANBgkqhkiG9w0B
AQsFAAOCAQEAE/RuzCheLIbHbrurGet6eaVx9DPYydNiKVBSAKO+5iuR84mQzhoT
nx5CN109xu5ml5baCYZsSnn6D7usC092bPA/kRCGxt29gkjpWA74tJHqIhVWgbM
mOrLiShoelewv+wRl0oVRChlTfKtXO68TUK6vrqpw76hKfOHIA7b2h1IIMdq6VA/
+A5FQ0xqYfGKdVd2RaINpzI8mqZiszqw+7E6j1PL5k4tftWEaYpfGPLVesFEyJEL
gHBUIPt8TIbaMYI8qUqMB/emnLXeKQ5PRxdRnleA3h8jfq3D1CQRTLjmdL3tpFwg
qopM6et5ZKqShX4T87BsgZIoiquzXqsuHg==
-----END CERTIFICATE-----"
  </pkITP>
</cloudCertStore>
</fvTenant>
</polUni>

```

Creating an HTTPS Service Policy Using the REST API

This section demonstrates how to create an HTTPS service policy using the REST API.



Note A listener can have multiple certificates. The certificate options are:

- ELBSecurityPolicy-2016-08 – The default when no security policy is chosen.
- ELBSecurityPolicy-FS-2018-06
- ELBSecurityPolicy-TLS-1-2-2017-01
- ELBSecurityPolicy-TLS-1-2-Ext-2018-06
- ELBSecurityPolicy-TLS-1-1-2017-01
- ELBSecurityPolicy-2015-05
- ELBSecurityPolicy-TLS-1-0-2015-04

If you use multiple certificates, you must specify the default certificate. The default is specified using the **defaultCert** property in **cloudRsListenerToCert**.

Before you begin

You have already configured a key ring certificate.

To create an HTTPS service policy:

```

<polUni>
  <fvTenant name="t2">
    <vnsAbsGraph name="CloudGraph" type="cloud" status="">
      <vnsAbsNode funcType="GoTo" name="N1" managed="yes">
        <cloudSvcPolicy tenantName="t2" contractName="httpFamily" subjectName="consubj">
          <cloudListener name="https_listener" port="443" protocol="https"
secPolicy="eLBSecurityPolicy-2016-08" status="">
            <cloudRsListenerToCert defaultCert="yes" certStore="iam"
tDn="uni/tn-t2/certstore/keyring-lbCert" status=""/>
              <cloudListenerRule name="defaultRule" default="yes" priority="100" status="">
                <cloudRuleAction type="forward" port="80" protocol="http"
epgdn="uni/tn-t1/cloudapp-ap/cloudepg-ep1">
                  </cloudRuleAction>
                </cloudListenerRule>
              </cloudListenerRule>
            </cloudListener>
          </cloudSvcPolicy>
        </vnsAbsNode>
      </vnsAbsGraph>
    </fvTenant>
  </polUni>

```



CHAPTER 7

Cisco Cloud APIC Statistics

- [About Cisco Cloud APIC Statistics, on page 153](#)
- [AWS Networking Interface Statistics Collection, on page 153](#)
- [Cisco Cloud APIC Endpoints and cloudEPg Statistics Processing, on page 154](#)
- [Cisco Cloud APIC Statistics Filters, on page 154](#)
- [AWS Transit Gateway Statistics, on page 155](#)
- [Enabling VPC Flow Logs, on page 156](#)
- [Cloud Router Statistics, on page 159](#)

About Cisco Cloud APIC Statistics

Cisco Cloud Application Policy Infrastructure Controller (APIC) supports statistics that are collected from the cloud routers. Additionally, it supports statistics that are derived by processing Amazon Web Services (AWS) flow logs. Because AWS flow logs is not a free service, the Cisco Cloud APIC provides a policy that allows you to control this feature. This feature is not enabled by default.

For more information about CloudWatch and flow logs, see "VPC Flow Logs" for Amazon Virtual Private Cloud on the AWS website.

Beginning in Cisco Cloud APIC Release 5.0(1), you can do the following:

- You can use filters to see specific information from the AWS flow logs. You can define up to eight filters for a given flow log policy (or VPC) at the same time. You can filter for a combination of source or destination IP address, port and protocol. See [Cisco Cloud APIC Statistics Filters, on page 154](#) for more information.
- You can collect statistics for traffic to and from AWS Transit Gateways. See the section [AWS Transit Gateway Statistics, on page 155](#) in this guide.

AWS Networking Interface Statistics Collection

AWS provides the nonreal-time IP traffic information per network interface through flow logs. Cisco Cloud APIC provides a policy for enabling flow logs per `cloudCtxProfile`. Because the `cloudCtxProfile` maps to a VPC in AWS, enabling flow logs per `cloudCtxProfile` or VPC means that you enabled flow logs for each interface belonging to that VPC. Once flow logs are enabled, flow records are periodically pushed to AWS Cloudwatch. The Cisco Cloud APIC then periodically polls AWS CloudWatch for these flow records

and parses these records to extract statistics. Because it can take up to 15 minutes to publish flow records to CloudWatch, the Cisco Cloud APIC delays its flow logs query to CloudWatch by 15 minutes too. This means that there is a lag between the flow logs being present in CloudWatch and the corresponding statistics showing up on the Cisco Cloud APIC. Cisco Cloud APIC does not process flow records that take longer than 15 minutes to publish to CloudWatch.

Cisco Cloud APIC Endpoints and cloudEPg Statistics Processing

The Cisco Cloud APIC extracts the following statistics for each AWS networking endpoint that has flow logs present in CloudWatch:

- Number of bytes or packets sent (egress)
- Number of bytes or packets received (ingress)
- Number of bytes or packets rejected (egress drop)
- Number of bytes or packets dropped (ingress drop)

These statistics are associated with the `cloudEpInfoHolder` observable.

Also, the Cisco Cloud APIC maps the flow log records to one or more per region `cloudEPg` objects. This is because a `cloudEPg` can be present in multiple regions. These statistics are associated with the `cloudRgInfoHolder` observable. This observable is a child of `cloudEPg` and the accumulation of statistics for the `cloudRgInfoHolder` children results in statistics for `cloudEPg`. The `cloudEPg` supports the following statistics:

- Number of bytes or packets sent (egress)
- Number of bytes or packets received (ingress)
- Number of bytes or packets rejected (egress drop)
- Number of bytes or packets dropped (ingress drop)

The `cloudEPg` statistics are aggregated up `fvApp` and then up `fvTenant`.

Cisco Cloud APIC Statistics Filters

Beginning in Cisco Cloud Application Policy Infrastructure Controller Release 5.0(1), you can use filters to see specific information from the Amazon Web Services (AWS) flow logs.

Statistics are collected for each endpoint on which the filter is deployed. The filters enable you to see information about a flow, filtered by a combination of source or destination IP address, port, and protocol. You can define up to eight filters for a given AWS log group at the same time.

A statistics filter has the following three attributes:

- **PeerIP:** The IPv4 address to filter
- **PeerPort:** The port number to listen to
- **Protocol:** The protocol number to listen to



Note We recommend that you configure statistics filters using the Cisco Cloud APIC GUI. You can alternatively use REST API; however, if you do and then switch to the GUI, the feature will appear incomplete. You should stick to the method that you choose.

Use of statistics filters depend on enabling Virtual Private Cloud (VPC) flow log; you must enable the logs before you configure the statistics filters.

Flow logs, which are stored in AWS CloudWatch, consist of flow log records. Cisco Cloud Application Policy Infrastructure Controller (APIC) extracts statistics by parsing the flow log records.

It can take up to 15 minutes from the occurrence of a particular flow record to its being present in AWS CloudWatch. Cisco Cloud APIC polls for flow records that occurred 15 minutes or more in the past. It does not process flow records that take longer than 15 minutes to appear in AWS CloudWatch.

AWS Transit Gateway Statistics

You can collect statistics for traffic going through Amazon Web Services (AWS) Transit Gateways on both the infra tenant and the user tenant. Statistics reported for user tenant represent the traffic of an attachment between an user VPC and an AWS Transit Gateway. Statistics reported from infra tenant represents the traffic of an attachment between an infra VPC and a Transit Gateway.

The following statistics are collected for AWS Transit Gateway:

- Ingress packets
- Ingress packet bytes
- Ingress packet drops
- Ingress packet drop bytes
- Egress packets
- Egress packet bytes
- Egress packet drops
- Egress packet drop bytes

You can enable infra tenant Transit Gateway statistics collection from the Cisco Cloud Application Policy Infrastructure Controller **Setup - Region Management** page. See the section "Set Up the Cloud Site to Use AWS Transit Gateway" in [Increasing Bandwidth Between VPCs by Using AWS Transit Gateway](#).

You can enable user tenant Transit Gateway statistics collection by enabling flow logs on the user VPC. See the sections [Enabling VPC Flow Logs, on page 156](#) and [Enabling VPC Flow Logs Using the Cisco Cloud APIC GUI, on page 156](#) in this guide.

To view AWS Transit Gateway statistics, in the Cisco Cloud APIC GUI, click the **Statistics** tab and then click **AWS Transit Gateway** in the left navigation pane. The central pane displays the information.

Enabling VPC Flow Logs

Steps to enable VPC Flow Logs:

1. Define a log group policy.
2. Define a flow log policy and associate the log group that you defined in the first step.
3. Associate the flow log policy to one or more `cloudCtxProfile`.

Log group properties:

- **name**—The location in CloudWatch where flow logs are sent.



Note The actual log group name that is programmed in AWS is the concatenation of `<tenant name><cloudCtxProfile name><log group name>`.

- **retention**—The length of duration for storing the logs in CloudWatch. The default is 5-days.

Flow log properties:

- **trafficType**—The type of traffic to collect. Supported types are **all**, **accepted only**, and **rejected only**. The default is **all**.

Enabling VPC Flow Logs Using the Cisco Cloud APIC GUI

This section explains how to enable VPC flow logs using the Cisco Cloud APIC GUI.



Note If you want to use filters to see specific information from AWS flow logs, perform the optional steps in this procedure.

Step 1 Click the **Navigation** menu and choose **Application Management > Tenants**.

The **Tenants** window appears with the tenants listed as rows in a summary table.

Step 2 Double-click on a tenant.

The tenant dialog box appears over the Work pane. The tenant dialog box displays the **Overview**, **Cloud Resources**, **Application Management**, **Statistics**, and **Event Analytics** tabs.

Step 3 Click the **Statistics** tab.

The **EPGs**, **CCRs**, and **Flow Log Collection** subtabs appear.

Step 4 Click **Flow Log Collection**.

The **Flow Log Collection Settings** information appears at the top of the dialog box with the **edit** icon in the top-right corner.

Step 5 Click the **edit** icon.

The **Flow Log Collection Settings** dialog box appears.

Step 6 Enter the appropriate values in each field as listed in the following *Flow Log Collection Settings Dialog Box Fields* table then continue.

Table 38: Flow Log Collection Settings Dialog Box Fields

Properties	Description
Type of Traffic to be Logged	Click the Type of Traffic to be Logged drop-down list and choose one of the following options: <ul style="list-style-type: none"> • All Traffic (default) • Accepted Only Traffic • Rejected Only Traffic
Destination	Click the Destination drop-down list and choose CloudWatch (default).
Retention	Click the Retention drop-down list and chose from the following options: <ul style="list-style-type: none"> • 1 day • 3 days • 5 days (default) • 1 month • 13 months • 18 months • 2 months • 3 months • 4 months • 5 months • 6 months • 1 week • 2 weeks • 1 year • 10 years • 2 years • 5 years

Step 7 (Optional) Add flow filters to get information about source and destination IP addresses, ports, or protocols by completing the following tasks:

For information about statistics filters, see the section [Cisco Cloud APIC Statistics Filters, on page 154](#).

a) Click **Add Flow Filters** at the bottom of the **Flow Log Collection Settings** dialog box.

Fields for the filter attributes appear.

After you click on the **Add Flow Filters** button, you will see a new filter being created; fill out the attributes.

b) In the **Peer IP** field, enter the IPv4 IP address of the peer.

The address needs to be in the format x.x.x.x/x. It tells the filter which network to monitor. An address of 0.0.0.0/0 will match all.

c) (Optional) From the **Protocol** drop-down list, choose a protocol to listen to.

The choices are integers from 0 to 255. Entering 255 will match any protocol. Well-known protocols are translated when text format is given:

<ul style="list-style-type: none"> • "icmp": 1 • "igmp": 2 • "tcp": 6 • "egp": 8 	<ul style="list-style-type: none"> • "igmp": 9 • "l2tp": 115 • "udp": 17 • "icmpv6": 58 	<ul style="list-style-type: none"> • "eigrp": 88 • "ospfigp": 89 • "pim": 103
--	---	--

d) (Optional) In the **Peer Port** field, enter the port number to listen to.

This number must be an integer from 0 to 65535 or text input for a well-known port number. Entering 0 will match any port. Well-known protocols are translated when text format is given:

<ul style="list-style-type: none"> • "dns": 53 • "ftpData": 20 • "smtp": 25 	<ul style="list-style-type: none"> • "http": 80 • "https": 443 	<ul style="list-style-type: none"> • "rtsp": 554 • "pop3": 110
--	--	--

e) (Optional) Check the **Active** check box and then click the check icon.

Step 8 Click **Save**.

Enabling VPC Flow Logs Using the REST API

This section demonstrates how to enable VPC flow logs using the REST API.

Step 1 Create a log group:

```
<cloudAwsLogGroup name="lg1" retention="days-3" status="">
  </cloudAwsLogGroup>
```

Step 2 Create a flow log policy:


```
<cloudAwsFlowLogPol name="flowLog1" trafficType="ALL" status="">
  <cloudRsToLogGrp tDn="uni/tn-t20/loggrp-lg1" status=""/>
</cloudAwsFlowLogPol>
```

Step 3 Create a relationship from a CtxProfile to a flow log policy:

```
<cloudCtxProfile name=" vrf1" status="">
  <cloudRsCtxToFlowLog tnCloudAwsFlowLogPolName="flowLog1" status=""/>
</cloudCtxProfile>
```

Cloud Router Statistics

These statistics are available for the cloud router:

- Ingress packets
- Egress packets
- Ingress bytes
- Egress bytes

The Cisco Cloud Application Policy Infrastructure Controller (APIC) collects and stores the cloud router statistics by the following granularities:

- 15-minutes
- 1-hour
- 1-month
- 1-year

Collection Mechanism

Each cloud router instance captures and stores the previously mentioned 4-stat values for each physical and tunnel interface.

The Cisco Cloud Application Policy Infrastructure Controller (APIC) queries the cloud routers for these statistics and maps the response to cloud router statistics on the Cisco Cloud APIC. The statistics query repeats every 5 minutes for as long as the tunnel is up and operational.

Raw Statistics


The raw statistics are stored under 2 Dns:

- uni/tn-*<infraTenant>*/ctx-*<infraCtx>*/region-*<infraRegion>*/router-*<csrname>*/to-*<ip or user-region>*/tunn-*<tunnel-id>*
- uni/tn-*<userTenant>*/ctx-*<userCtx>*/region-*<userRegion>*/region-*<infraRegion>*/router-*<csrname>*/tunn-*<tunnel-id>*

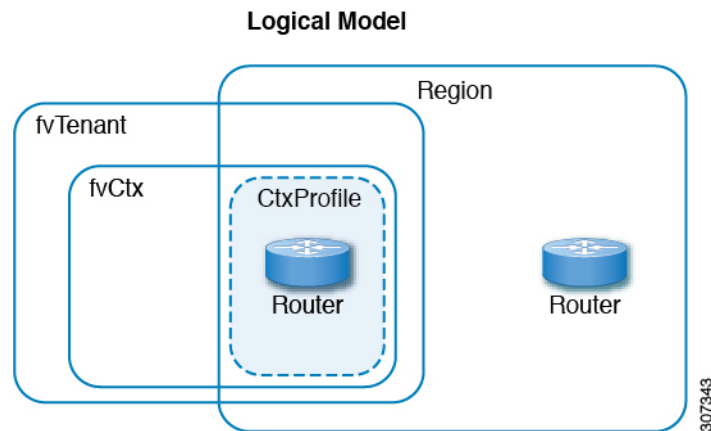
**Note**

- The second Dn holder is the statistics as seen from the user endpoints connected to the cloud router. These statistics are hence flipped (Ingress on the CCR becomes egress on the user region)
- Not all tunnels have a corresponding user dn. This is only applicable to internal tunnels. External tunnels statistics are only available on the 1st Dn.

In the following figure, internal tunnels are between the user VPC and infra VPC. The infra VPC contains the CCR router. The user VPC can contain the CCR or VGW router. Cisco Cloud APIC creates these tunnels. As a result, statistics are available for both the infra side and the user side. External tunnels are between infra VPC and an external IP address. Statistics are only available on the infra side (Dn-1).



In the logical model diagram, a tenant can be infra or a user tenant. You configure a VRF (or `fvCtx`) to be inside a tenant (per tenant). A VRF can be within one region or span across multiple regions.



Aggregated Statistics

Statistics are aggregated at each parent level of the DN. For the preceding case, statistics on tunnel, statistics are aggregated on to the destination IP, cloud router, region, vrf (ctx), and tenant.

For example, if you want to find the egress packets from the infra cloud router to a user region, it is available under `uni/tn-<infraTenant>/ctx-<infraCtx>/region-<infraRegion>/router-<csrname>/to-<ip or user-region>/`

If you want to get all the packets between user region1 and infra region2, it is available under `uni/tn-<userTenant>/ctx-<userCtx>/region-<userRegion>/region-<infraRegion>/`

Also, if you want to find statistics per `cloudCtxProfile`, it is available under `uni/tn-<userTenant>/ctx-<userCtx>/region-<userRegion>/` OR `uni/tn-<infraTenant>/ctx-<infraCtx>/region-<infraRegion>/`.

Cloud Router GUI Statistics

In the Cisco Cloud APIC GUI, statistics are visible available under the tenant, VRF, infra region, and cloud context profile.

For Amazon Web Services (AWS) Transit Gateway stats, in the **Cloud Context Profile** work pane, click **AWS Transit Gateway**. For all other stats, in the **Cloud Context Profile** work pane, click **Endpoints**.



CHAPTER 8

Cisco Cloud APIC Security

This chapter contains the following sections:

- [Access, Authentication, and Accounting, on page 163](#)
- [Configuring TACACS+, RADIUS, LDAP and SAML Access, on page 164](#)
- [Configuring HTTPS Access, on page 171](#)

Access, Authentication, and Accounting

Cisco Cloud Application Policy Infrastructure Controller (Cloud APIC) policies manage the authentication, authorization, and accounting (AAA) functions. 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 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.

For more access, authentication, and accounting configuration information, see *Cisco APIC Security Configuration Guide, Release 4.0(1)* at <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html>.

Configuration

The admin account is configured in the initial configuration script, and the admin is the only user when the system starts.

Configuring a Local User

Refer to [Creating a Local User Using the Cisco Cloud APIC GUI, on page 108](#) to configure a Local User and associate it to the OTP, SSH Public Key, and X.509 User Certificate using the Cisco Cloud APIC GUI.

Configuring TACACS+, RADIUS, LDAP and SAML Access

The following topics describe how to configure TACACS+, RADIUS, LDAP and SAML access for the Cisco Cloud APIC.

Overview

This topic provides step-by-step instructions on how to enable access to the Cisco Cloud APIC for RADIUS, TACACS+, LDAP, and SAML users, including ADFS, Okta, and PingID.

For additional TACACS+, RADIUS, LDAP, and SAML information, see *Cisco APIC Security Configuration Guide, Release 4.0(1)* at <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html>.

Configuring Cloud APIC for TACACS+ Access

Before you begin

- The Cloud Application Policy Infrastructure Controller (Cloud APIC) is online.
- The TACACS+ server host name or IP address, port, and key are available.
- The Cloud APIC management endpoint group is available.

Step 1

In the Cloud APIC, create the **TACACS+ Provider**.

- a) Click the **Global Create** icon.

The **Global Create** menu appears.

- b) Scroll down until you see the **Administrative** area, then click **Create Provider** under the **Administrative** area.

The **Create Provider** dialog box appears.

- c) In the **Host Name/IP Address** field, enter the Host Name/IP Address of the provider.
- d) In the **Description** field, enter a description of the provider.
- e) Click the **Type** drop-down list and choose **TACACS+**.
- f) In **Settings** section, specify the **Key**, **Port**, **Authentication Protocol**, **Timeout**, **Retries**, **Management EPG**. Select either **Enabled** or **Disabled** for **Server Monitoring**.

Step 2

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

- a) Click the **Global Create** icon.

The **Global Create** menu appears.

- b) Click the drop-down arrow below the **Global Create** search box and choose **Administrative**.

A list of **Administrative** options appear in the **Global Create** menu.

- c) From the **Administrative** list in the **Global Create** menu, click **Create Login Domain**.

The **Create Login Domain** dialog box appears.

- d) Enter the appropriate values in each field as listed in the following *Create Login Domain Dialog Box Fields* table then continue.

Properties	Description
General	
Name	Enter the name of the Login Domain
Description	Enter the description of the Login Domain.
Settings	
Realm	Choose TACACS+ from the dropdown menu
Providers	<p>To choose a Provider(s):</p> <ol style="list-style-type: none"> 1. Click Add Providers. The Select Providers dialog appears. 2. Click to choose a provider(s) in the column on the left. 3. Click Select. You return to the Create Login Domain dialog box.

- e) Click **Save** to save the configuration.

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.

Configuring Cloud APIC for RADIUS Access

Before you begin

- The Cloud Application Policy Infrastructure Controller (Cloud APIC) is online.
- The RADIUS server host name or IP address, port, and key are available.
- The Cloud APIC management endpoint group is available.

Step 1 In the Cloud APIC, create the **RADIUS Provider**.

- Click the **Global Create** icon.
The **Global Create** menu appears.
- Scroll down until you see the **Administrative** area, then click **Create Provider** under the **Administrative** area.
The **Create Provider** dialog box appears.

- c) In the **Host Name/IP Address** field, enter the Host Name/IP Address of the provider.
- d) In the **Description** field, enter a description of the provider.
- e) Click the **Type** drop-down list and choose **RADIUS**.
- f) In the **Settings** section, specify the **Key, Port, Authentication Protocol, Timeout, Retries, Management EPG**. Select either **Enabled** or **Disabled** for **Server Monitoring**.

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

- a) Click the **Global Create** icon.

The **Global Create** menu appears.

- b) Click the drop-down arrow below the **Global Create** search box and choose **Administrative**

A list of **Administrative** options appear in the **Global Create** menu.

- c) From the **Administrative** list in the **Global Create** menu, click **Create Login Domain**.

The **Create Login Domain** dialog box appears.

- d) Enter the appropriate values in each field as listed in the following *Create Login Domain Dialog Box Fields* table then continue.

Properties	Description
General	
Name	Enter the name of the Login Domain
Description	Enter the description of the Login Domain.
Settings	
Realm	Choose RADIUS from the dropdown menu
Providers	To choose a Provider(s): <ol style="list-style-type: none"> 1. Click Add Providers. The Select Providers dialog appears. 2. Click to choose a provider(s) in the column on the left. 3. Click Select. You return to the Create Login Domain dialog box.

- e) Click **Save** to save the configuration.

What to do next

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

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

Refer to the section *Configuring a Cisco Secure Access Control Server for RADIUS and TACACS+ Access to the APIC* in the *Cisco APIC Security Configuration Guide, Release 4.0(1)* at <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html>.

Configuring LDAP Access

There are two options for LDAP configurations:

- Configure a Cisco AVPair
- Configure LDAP group maps in the cloud APIC

The following sections contain instructions for both configuration options.

Configuring Windows Server 2008 LDAP for APIC Access with Cisco AVPair

Refer to the section *Configuring Windows Server 2008 LDAP for APIC Access with Cisco AVPair* in the *Cisco APIC Security Configuration Guide, Release 4.0(1)* at <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html>.

Configuring Cloud APIC for LDAP Access

Before you begin

- The Cloud Application Policy Infrastructure Controller (Cloud APIC) is online.
- The LDAP server host name or IP address, port, bind DN, Base DN, and password are available.
- The cloud APIC management endpoint group is available.

Step 1

In the Cloud APIC, create the **LDAP Provider**.

- a) On the menu bar, choose **Administrative > Authentication**.
- b) In the Work pane, click on **Providers** tab and then click on the **Actions** drop-down and select **Create Provider**.
- c) In the **Host Name/IP Address** field, enter the Host Name/IP Address of the provider.
- d) In the **Description** field, enter a description of the provider.
- e) Click the **Type** drop-down list and choose **LDAP**.
- f) Specify the **Bind DN**, **Base DN**, **Password**, **Port**, **Attribute**, **Filter Type** and **Management EPG**.

- Note**
- The bind DN is the string that the Cloud APIC uses to log in to the LDAP server. The Cloud 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 Cloud APIC searches for the remote user account. This is where the password is validated. Filter is used to locate the attribute that the Cloud APIC requests to use for the *cisco-av-pair*. This contains the user authorization and assigned RBAC roles for use on the Cloud APIC. The Cloud 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**.

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

- On the menu bar, choose **Administrative > Authentication**.
- In the Work pane, click on **Login Domains** tab and then click on the **Actions** drop-down and select **Create Login Domain**.
- Enter the appropriate values in each field as listed in the following *Create Login Domain Dialog Box Fields* table then continue.

Properties	Description
General	
Name	Enter the name of the Login Domain
Description	Enter the description of the Login Domain.
Settings	
Realm	Choose LDAP from the dropdown menu
Providers	<p>To choose a Provider(s):</p> <ol style="list-style-type: none"> 1. Click Add Providers. The Select Providers dialog appears. 2. Click to choose a provider(s) in the column on the left. 3. Click Select. You return to the Create Login Domain dialog box.

Properties	Description
Authentication Type	<ol style="list-style-type: none"> 1. Select Cisco AV Pairs, if provider(s) was configured with CiscoAVPair as the Attribute. 2. Select LDAP Group Map Rules, if provider(s) was configured with memberOf as the Attribute. <ol style="list-style-type: none"> a. Click Add LDAP Group Map Rule. The dialog box appears. b. Specify the map rule Name, Description (optional), and Group DN. c. Click the + next to Add Security Domain. The dialog box appears. d. Click the + to access the Role name and Role Privilege Type (Read or Write) fields. Click check mark. e. Repeat step 4 to add more roles. Then click Add. f. Repeat step 3 to add more security domains. Then click Add.

d) Click **Save** on Create Login Domain dialog box.

Configuring Cloud APIC for SAML Access

The following sections provide detailed information on configuring Cloud APIC for SAML access.

About SAML

Refer to the section *About SAML* in the *Cisco APIC Security Configuration Guide, Release 4.0(1)* at <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html>.

Basic Elements of SAML

Refer to the section *Basic Elements of SAML* in the *Cisco APIC Security Configuration Guide, Release 4.0(1)* at <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html>.

Supported IdPs and SAML Components

Refer to the section *Supported IdPs and SAML Components* in the *Cisco APIC Security Configuration Guide, Release 4.0(1)* at <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html>.

Configuring Cloud APIC for SAML Access



Note SAML based Authentication is only for Cloud APIC GUI and not for REST.

Before you begin

- The SAML server host name or IP address, and the IdP's metadata URL are available.
- The Cloud APIC management endpoint group is available.
- Set up the following:
 - Time Synchronization and NTP
 - Configuring a DNS Provider Using the GUI
 - Configuring a Custom Certificate for Cisco ACI HTTPS Access Using the GUI

Step 1

In the Cloud APIC, create the **SAML Provider**.

- a) On the menu bar, choose **Administrative > Authentication**.
- b) In the **Work** pane, click on **Providers** tab and then click on the **Actions** drop-down and select **Create Provider**.
- c) In the **Host Name/IP Address** field, enter the Host Name/IP Address of the provider.
- d) In the **Description** field, enter a description of the provider.
- e) Click the **Type** drop-down list and choose **SAML**.
- f) In **Settings** pane, perform following:
 - Specify the IdP metadata URL:
 - In case of AD FS, IdP Metadata URL is of the format *https://<FQDN ofADFS>/FederationMetadata/2007-06/FederationMetadata.xml*.
 - In case of Okta, to get the IdP Metadata URL, copy the link for **Identity Provider Metadata** in the **Sign On** section of the corresponding SAML Application from the Okta server.
 - Specify the **Entity ID** for the SAML-based service.
 - Configure the **HTTPS Proxy for Metadata URL** if it is needed to access the IdP metadata URL.
 - Select the **Certificate Authority** if IdP is signed by a Private CA.
 - Select the **Signature Algorithm Authentication User Requests** from the drop-down.
 - Select checkbox to enable **Sign SAML Authentication Requests**, **Sign SAML Response Message**, **Sign Assertions in SAML Response**, **Encrypt SAML Assertions**.
- g) Click **Save** to save the configuration.

Step 2

Create the login domain for SAML.

- a) On the menu bar, choose **Administrative > Authentication**.
- b) In the **Work** pane, click on the **Login Domains** tab and then click on the **Actions** drop-down and select **Create Login Domain**.

- c) Enter the appropriate values in each field as listed in the following Create Login Domain Dialog Box Fields table then continue.

Properties	Description
General	
Name	Enter the name of the Login Domain
Description	Enter the description of the Login Domain.
Settings	
Realm	Choose SAML from the dropdown menu
Providers	To choose a Provider(s): <ol style="list-style-type: none"> 1. Click Add Providers. The Select Providers dialog appears. 2. Click to choose a provider(s) in the column on the left. 3. Click Select. You return to the Create Login Domain dialog box.

- d) Click **Save** to save the configuration.

Setting Up a SAML Application in Okta

Refer to the section *Setting Up a SAML Application in Okta* of *Cisco APIC Security Configuration Guide, Release 4.0(1)* at <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html>.

Setting Up a Relying Party Trust in AD FS

Refer to the section *Setting Up a Relying Party Trust in AD FS* in the *Cisco APIC Security Configuration Guide, Release 4.0(1)* at <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html>.

Configuring HTTPS Access

The following sections describe how to configure HTTPS access.

About HTTPS Access

This article provides an example of how to configure a custom certificate for HTTPS access when using Cisco ACI.

For more information, see the section *HTTPS Access* in the *Cisco APIC Security Configuration Guide, Release 4.0(1)* at <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html>.

Guidelines for Configuring Custom Certificates

- Wild card certificates (such as *.cisco.com, which is used across multiple devices) and its associated private key generated elsewhere are not supported on the Cisco Cloud APIC as there is no support to input the private key or password in the Cisco Cloud APIC. Also, exporting private keys for any certificates, including wild card certificates, is not supported.
- 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 Cloud 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 Cloud 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.
- Only one Certificate Based Root can be active per pod.
- Client Certificate based authentication is not supported for this release.

Configuring a Custom Certificate for Cisco ACI HTTPS Access Using the GUI

Determine from which authority you will obtain the trusted certification so that you can create the appropriate Certificate Authority.

Before you begin

CAUTION: PERFORM THIS TASK ONLY DURING A MAINTENANCE WINDOW AS THERE IS A POTENTIAL FOR DOWNTIME. Expect a restart of all web servers on Cloud APIC during this operation.

Step 1 On the menu bar, choose **Administrative > Security**.

- Step 2** In the Work pane, click on **Certificate Authorities** tab and then click on the **Actions** drop-down and select **Create Certificate Authority**.
- Step 3** In the **Create Certificate Authority** dialog box, in the **Name** field, enter a name for the certificate authority and in the **Description** field, enter a description.
- Step 4** Select **System** in the **Used for** field.
- 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 Cloud Application Policy Infrastructure Controller (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 **Save**.
- Step 7** On the menu bar, choose **Administrative > Security**.
- Step 8** In the Work pane, click on the **Key Rings** tab, then click on the **Actions** drop-down and select **Create Key Ring**.
- Step 9** In the **Create Key Ring** dialog box, in the **Name** field, enter a name for the certificate authority and in **Description** enter description.
- Step 10** Select **System** in the **Used for** field.
- Step 11** For the **Certificate Authority** field, click on **Select Certificate Authority** and select the Certificate Authority that you created earlier.
- Step 12** Select either **Generate New Key** or **Import Existing Key** for the field **Private Key**. If you select **Import Existing Key**, enter a private key in the **Private Key** text box.
- Step 13** Select modulus from the **Modulus** drop-down. menu
- Step 14** In the **Certificate** field, do not add any content.
- Step 15** Click **Save**.
- In the **Work** pane, in the **Key Rings** area, the **Admin State** for the key ring created displays **Started**.
- Step 16** Double-click on the created Key Ring to open **Key Ring** *key\_ring\_name* dialog box from the **Work** pane.
- Step 17** In the **Work** pane, click on **Create Certificate Request**.
- Step 18** In the **Subject** field, enter the fully qualified domain name (FQDN) of the Cloud APIC.
- Step 19** Fill in the remaining fields as appropriate.
- Step 20** Click **Save**.
- The **Key Ring** *key\_ring\_name* dialog box appears.
- Step 21** Copy the contents from the field Request to submit to the **Certificate Authority** for signing.
- Step 22** From the **Key Ring** *key\_ring\_name* dialog box, click on edit icon to display the **Key Ring** *key\_ring\_name* dialog box.
- Step 23** In the **Certificate** field, paste the signed certificate that you received from the certificate authority.
- Step 24** Click **Save** to return to the **Key Rings** work pane.
- The key is verified, and in the **Work** pane, the **Admin State** changes to **Completed** and is now ready for use in the HTTPS policy.
- Step 25** Navigate to **Infrastructure > System Configuration**, then click the **Management Access** tab.

**Step 26** Click the edit icon on the **HTTPS** work pane to display the **HTTPS Settings** dialog box.

**Step 27** Click on **Admin Key Ring** and associate the Key Ring that you created earlier.

**Step 28** Click **Save**.

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 Cloud APIC.





## CHAPTER 9

# Configuration Drifts

---

- [Configuration Drift Notifications and Faults, on page 175](#)
- [Accessing the Main Configuration Drift Page, on page 176](#)
- [Checking for Missing Contracts Configuration, on page 179](#)
- [Checking for Missing EPGs Configuration, on page 180](#)
- [Checking for Missing VRFs Configuration, on page 182](#)
- [Configuration Drift Troubleshooting, on page 183](#)

## Configuration Drift Notifications and Faults

When you deploy Cisco ACI in a public cloud, you will perform most of the fabric configuration from the Cloud APIC. However, there may be cases where you or another cloud administrator changes the deployed configuration directly in the cloud provider's GUI using the tools provided by AWS or Azure. In these cases, the intended configuration you deployed from the Cloud APIC and the actual configuration in the cloud site may become out of sync, we call this a configuration drift.

Starting with release 5.0(2), Cloud APIC provides visibility into any security policy (contracts) configuration discrepancy between what you deploy from the Cloud APIC and what is actually configured in the cloud site.



---

**Note**

- Beginning with release 25.0(1), configuration drift information is available for EPGs and VRFs, in addition to contracts.
- Beginning with release 25.0(4), contract drift information is available for contracts with or without Layer 4 to Layer 7 service graphs attached.

See [Updates in Release 25.0\(4\), on page 176](#) for more information.

---

There are two aspects to analyzing configuration drift:

- Have all the fabric elements configured in the Cloud APIC and intended to be deployed in the cloud fabric been properly deployed?

This scenario can occur due to user configuration errors in Cloud APIC that could not be deployed in the cloud, connection or API issues on the cloud provider end, or if a cloud administrator manually deletes or modifies security rules directly in the cloud provider's UI. Any intended but missing configurations may present an issue for the Cloud APIC fabric.

- Are there any additional configurations that exist in the cloud but were not intended to be deployed from the Cloud APIC?

Similarly to the previous scenario, this can occur if there are connection or API issues or if a cloud administrator manually creates additional security rules directly in the cloud provider's UI. Any existing but not intended configuration may present issues.

### Updates in Release 25.0(4)

Beginning with release 25.0(1), configuration drift information is available for EPGs and VRFs, in addition to contracts.

Beginning with release 25.0(4), the following changes have been made for configuration drift:

- Configuration drift is now enabled by default.
- Prior to release 25.0(4), configuration drift information was not available for contracts that had Layer 4 to Layer 7 service graphs attached. Beginning with release 25.0(4), contract drift information is now available for contracts with or without Layer 4 to Layer 7 service graphs attached. See [Deploying Layer 4 to Layer 7 Services, on page 135](#) for more information.
- Configuration drift information is now consolidated under a single page, located at **Cloud Resources > Drifts**.

| Object                                                                                                                                                                                             | Status    | Drift Type   | Last Configuration Update     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|--------------|-------------------------------|
| /subscriptions/630daf13-e0e6-40bf-b87e-17844309a7f3/resourceGroups/CAPIC_tnt_vrfl1_eastus/providers/Microsoft.Network/networkSecurityGroups/ALB2-subnet/securityRules/AzureInfraBackendHealthPorts | Unmanaged | Extra Object | Jul 08 2022 01:43:56am -07:00 |
| /subscriptions/630daf13-e0e6-40bf-b87e-17844309a7f3/resourceGroups/CAPIC_tnt_vrfl1_eastus/providers/Microsoft.Network/networkSecurityGroups/ALB2-subnet/securityRules/AzureLoadBalancer            | Unmanaged | Extra Object | Jul 08 2022 01:32:32am -07:00 |
| /subscriptions/630daf13-e0e6-40bf-b87e-17844309a7f3/resourceGroups/CAPIC_tnt_vrfl1_eastus/providers/Microsoft.Network/networkSecurityGroups/vrf1-east-VM1-nsg/securityRules/SSH                    | Unmanaged | Extra Object | Jul 08 2022 01:32:32am -07:00 |

See [Accessing the Main Configuration Drift Page, on page 176](#) for more information.

## Accessing the Main Configuration Drift Page

Beginning with release 25.0(4), configuration drift information is now consolidated under a single **Drifts** page.

The **Drifts** page is used to provide the following pieces of information:

- To identify if something has been deleted
- To ensure that anything that should be present is correctly shown as present

**Step 1** Log in to your Cloud APIC GUI.

**Step 2** Navigate to the main configuration drift page:

**Cloud Resources > Drifts**

The consolidated **Drifts** page appears.

The screenshot shows the Cisco Cloud APIC GUI. On the left is a dark blue sidebar with navigation options: Dashboard, Topology, Cloud Resources (expanded), Regions, Virtual Networks, Routers, Security Groups, Endpoints, Virtual Machines, Network Services, Kubernetes Clusters, Native Services, **Drifts** (highlighted), Application Management, Operations, Infrastructure, and Administrative. The main content area is titled 'Drifts' and features a 'Detection Summary' section with the following data:

| Unmanaged Objects | Objects with Drifts | Last Drift Check              |
|-------------------|---------------------|-------------------------------|
| 32                | 0                   | Jul 08 2022 08:37:26am -07:00 |

Below the summary is a table of configuration drifts:

| Object                                                                                                                                                                                           | Status    | Drift Type   | Last Configuration Update     |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|--------------|-------------------------------|
| /subscriptions/630daf13-e0e6-40bf-b87e-17844309af73/resourceGroups/CAPIC_int_vr1_eastus/providers/Microsoft.Network/networkSecurityGroups/ALB2-subnet/securityRules/AzureInfraBackendHealthPorts | Unmanaged | Extra Object | Jul 08 2022 01:43:56am -07:00 |
| /subscriptions/630daf13-e0e6-40bf-b87e-17844309af73/resourceGroups/CAPIC_int_vr1_eastus/providers/Microsoft.Network/networkSecurityGroups/ALB2-subnet/securityRules/AzureLoadBalancer            | Unmanaged | Extra Object | Jul 08 2022 01:43:56am -07:00 |
| /subscriptions/630daf13-e0e6-40bf-b87e-17844309af73/resourceGroups/CAPIC_int_vr1_eastus/providers/Microsoft.Network/networkSecurityGroups/vr1-east-VM1-rsng/securityRules/SSH                    | Unmanaged | Extra Object | Jul 08 2022 01:32:32am -07:00 |
| /subscriptions/630daf13-e0e6-40bf-b87e-17844309af73/resourceGroups/CAPIC_int_vr1_eastus/providers/Microsoft.Network/networkSecurityGroups/vr1-east-VM1-rsng/securityRules/SSH                    | Unmanaged | Extra Object | Jul 08 2022 01:32:32am -07:00 |

In the **Drifts** page, you can see a summary of any configuration issues in your fabric.

The **Detection Summary** area provides an overview of how many configuration drifts were detected with managed or unmanaged objects, and the time when this information was last updated. If the inventory update timestamp is out of date, you can refresh the information by clicking the **Refresh** icon in the top right corner of this screen.

**Step 3** Use the information in the table below the **Detection Summary** area to find any configuration drifts.

- **Object:** Provides information on the object associated with the configuration drift.
- **Status:** Following are the different values that might appear in the **Status** column:
  - **Transient (low):** Drifts that are likely due to recent configuration changes. We recommend waiting for the fabric to stabilize and the drift will likely resolve on its own after the next configuration refresh.
  - **Presumed (medium):** Drifts that may or may not be transient. We recommend monitoring the status and troubleshoot the configuration should the drift persist.
  - **Raised (high):** Critical drifts. We recommend verifying the configuration on Cloud APIC and checking for any associated faults. Redeploying the configuration may help resolve communication issues between the Cloud APIC and cloud services. If the issue persists, check the tech-support logs.
  - **Unmanaged:** Configuration drifts related to extra inventory objects not created through the Cisco Cloud APIC.
- **Drift Type:** Following are the different values that might appear in the **Drift Type** column:
  - **Configuration:** External changes on the cloud providers site that could result in the intended configuration and the actual configuration being out of sync. Used for configuration drifts related to *EPGs* or *VRFs*.
  - **Rule:** External changes on the cloud providers site that could result in the intended security rules and the expected rules that are established through a contract being out of sync. Used for configuration drifts related to *contracts*.

- **Extra Object:** Used to show extra inventory objects that were not created through the Cisco Cloud APIC. Cisco Cloud APIC does not perform drift detection on these objects.

- **Last Configuration Update:** Provides information on when the last configuration update occurred.

**Step 4** Enter information in the filter line to filter the configuration drifts provided in the table, if necessary.

a) Click in the filter line below the **Detection Summary** area. The following filter types appear:

- Object
- Status
- Drift Type
- Last Configuration Update
- Parent Path

Select the appropriate type for your filter.

b) Click the necessary operator.

The options are:


- ==: The equal-to operator
- !=: The not-equal-to operator

c) Click the necessary drift type.

The options are `Extra Object`, `Rule`, and `Configuration`. See the explanations for the **Drift Type** field above for more information.

The entries in the table are filtered based on your selections above.

**Step 5** View additional information on a specific configuration drift, if necessary.

For any object listed in this page, you can bring up additional configuration drift information by clicking the appropriate line in the **Configuration Drifts** table. A side panel appears with additional information on this particular configuration drift; clicking the Details icon () automatically brings you to the appropriate **Cloud Mapping** page for this particular object.

Refer to the following sections for additional configuration drift information for specific objects:

- [Checking for Missing Contracts Configuration, on page 179](#)
- [Checking for Missing EPGs Configuration, on page 180](#)
- [Checking for Missing VRFs Configuration, on page 182](#)

# Checking for Missing Contracts Configuration

This section describes how to check for any contract settings you have configured from the Cloud APIC, but which have not been properly deployed to the cloud fabric.

- Step 1** Log in to your Cloud APIC GUI.
- Step 2** Click **Application Management > Contracts**.
- Step 3** Double-click the appropriate contract to bring up the **Overview** page for that contract.
- Step 4** Note the service graph information provided in the **Service Graph** area, if applicable.

Prior to release 25.0(4), configuration drift information was not available for contracts that had Layer 4 to Layer 7 service graphs attached. Beginning with release 25.0(4), contract drift information is now available for contracts with or without Layer 4 to Layer 7 service graphs attached. See [Deploying Layer 4 to Layer 7 Services, on page 135](#) for more information.

- Step 5** Click the **Cloud Mapping** tab.

The **Cloud Mapping** view displays all the information about the contract and the cloud resources it uses.

**Note** You can also navigate to this page by navigating to **Cloud Resources > Drifts**, then clicking the appropriate line in the **Configuration Drifts** table. A side panel appears with additional information on this particular configuration drift; clicking the Details icon (🔍) automatically brings you to the appropriate **Cloud Mapping** page for this particular object. See [Accessing the Main Configuration Drift Page, on page 176](#) for more information.

The screen is divided into four sections: **Detection Summary**, **Related Objects**, **Configuration Drifts**, and **Presented Cloud Resources**. Each section contains a table that lists the respective information about the contract you selected.

- The **Detection Summary** area provides an overview of how many configuration drifts were detected, number of intended and actual cloud resources configured, and the time when this information was last updated. If the inventory update timestamp is out of date, you can refresh the information by clicking the **Refresh** icon in the top right corner of this screen.

- The **Related Objects** area shows any other objects that have a relation to the contract, such as consumer and provider EPGs, and filters.
- The **Configuration Drifts** table lists all the issues with the contract rules. Specifically, all the contract rules that were intended to be deployed but are missing in the actual fabric configuration.

The table contains detailed information, such as the protocol used, port ranges, source and destination IP or group, consumer and provider EPGs, description of the issue, and the recommended action to resolve it. For each configuration drift, the **Status** field will indicate the severity and recommended action:

- **Transient (low)**: Drifts that are likely due to recent configuration changes. We recommend waiting for the fabric to stabilize and the drift will likely resolve on its own after the next configuration refresh.
  - **Presumed (medium)**: Drifts that may or may not be transient. We recommend monitoring the status and troubleshoot the configuration should the drift persist.
  - **Raised (high)**: Critical drifts. We recommend verifying the configuration on Cloud APIC and checking for any associated faults. Redeploying the configuration may help resolve communication issues between the Cloud APIC and cloud services. If the issue persists, check the tech-support logs.
- The **Presented Cloud Resources** table shows the information about all the resources that were properly configured in your cloud. This table is designed to provide you with better visibility into what rules are configured in your cloud for a specific contract.

---


## Checking for Missing EPGs Configuration

This section describes how to check for any EPG settings you have configured from the Cloud APIC, but which have not been properly deployed to the cloud fabric.

- 
- Step 1** Log in to your Cloud APIC GUI.
  - Step 2** Click **Application Management > EPGs**.
  - Step 3** Double-click the appropriate EPG to bring up the **Overview** page for that EPG.
  - Step 4** Click the **Cloud Mapping** tab.

The **Cloud Mapping** view displays all the information about the EPG and the cloud resources it uses.

The screenshot displays the Cisco Cloud APIC interface for EPG vrf1-epg1. The top navigation bar includes 'Overview', 'Topology', 'Cloud Resources', 'Application Management', 'Cloud Mapping', 'Statistics', and 'Event Analytics'. A notification banner states 'Detection of configuration drifts is still in beta.' Below this, the 'Detection Summary' section shows a table with columns: Configuration Drift Status (0 Drifts Found), Configured Cloud Resources (2), Expected Cloud Resources (2), and Last Drift Check (Jul 08 2022 08:46:00am -07:00). The 'Related Objects' section shows a table with columns: Application Security Group (2), Network Security Groups (2), Provider Contracts (2), and Consumer Contracts (0). The 'Configuration Drifts' section features a table with columns: Status, Logical DN, Cloud Provider ID, Drift Type, Description, and Recommendation. The table currently displays 'No rows found'. At the bottom, there is a 'Presented Cloud Resources' section.

**Note** You can also navigate to this page by navigating to **Cloud Resources > Drifts**, then clicking the appropriate line in the **Configuration Drifts** table. A side panel appears with additional information on this particular configuration drift; clicking the Details icon () automatically brings you to the appropriate **Cloud Mapping** page for this particular object. See [Accessing the Main Configuration Drift Page, on page 176](#) for more information.

The screen is divided into four sections: **Detection Summary**, **Related Objects**, **Configuration Drifts**, and **Presented Cloud Resources**. Each section contains a table that lists the respective information about the EPG you selected.

- The **Detection Summary** area provides an overview of how many configuration drifts were detected, number of intended and actual cloud resources configured, and the time when this information was last updated. If the inventory update timestamp is out of date, you can refresh the information by clicking the **Refresh** icon in the top right corner of this screen.
- The **Related Objects** area shows any other objects that have a relation to the EPG, such as security groups and contracts.
- The **Configuration Drifts** table lists all the issues with the security groups associated with the EPG. Specifically, all the security groups that were intended to be deployed but are missing in the actual fabric configuration.

The table contains detailed information, such as the logical DN, cloud provider ID, drift type, description of the issue, and the recommended action to resolve it. For each configuration drift, the **Status** field will indicate the severity and recommended action:

- **Transient (low)**: Drifts that are likely due to recent configuration changes. We recommend waiting for the fabric to stabilize and the drift will likely resolve on its own after the next configuration refresh.
- **Presumed (medium)**: Drifts that may or may not be transient. We recommend monitoring the status and troubleshoot the configuration should the drift persist.
- **Raised (high)**: Critical drifts. We recommend verifying the configuration on Cloud APIC and checking for any associated faults. Redeploying the configuration may help resolve communication issues between the Cloud APIC and cloud services. If the issue persists, check the tech-support logs.

- The **Presented Cloud Resources** table shows the information about all the resources that were properly configured in your cloud. This table is designed to provide you with better visibility into what security groups are associated with a specific EPG in your cloud.

## Checking for Missing VRFs Configuration

This section describes how to check for any VRF settings you have configured from the Cloud APIC, but which have not been properly deployed to the cloud fabric.

**Step 1** Log in to your Cloud APIC GUI.

**Step 2** Click **Application Management > VRFs**.

**Step 3** Double-click the appropriate VRF to bring up the **Overview** page for that VRF.

**Step 4** Click the **Cloud Mapping** tab.

The **Cloud Mapping** view displays all the information about the VRF and the cloud resources it uses.

**Note** You can also navigate to this page by navigating to **Cloud Resources > Drifts**, then clicking the appropriate line in the **Configuration Drifts** table. A side panel appears with additional information on this particular configuration drift; clicking the Details icon (🔍) automatically brings you to the appropriate **Cloud Mapping** page for this particular object. See [Accessing the Main Configuration Drift Page, on page 176](#) for more information.

The screen is divided into four sections: **Detection Summary**, **Related Objects**, **Configuration Drifts**, and **Presented Cloud Resources**. Each section contains a table that lists the respective information about the VRF you selected.

- The **Detection Summary** area provides an overview of how many configuration drifts were detected, number of intended and actual cloud resources configured, and the time when this information was last updated. If the inventory update timestamp is out of date, you can refresh the information by clicking the **Refresh** icon in the top right corner of this screen.



- The **Related Objects** area shows any other objects that have a relation to the VRF, such as security groups, CIDRs, and subnets.
- The **Configuration Drifts** table lists all the issues with the virtual networks, the CIDRs that are associated with the virtual networks, and the subnets within those CIDRs. Specifically, all the virtual networks, CIDRs, and subnets that were intended to be deployed but are missing in the actual fabric configuration.

Note that if there are configuration drifts at any one level, the table will show the configuration drift at that level and not any configuration drifts at the levels below it. For example, if a configuration drift occurs at a CIDR level and the corresponding subnets within that CIDR, the table will display the configuration drifts in the CIDR area but not the configuration drifts for the corresponding subnets within that CIDR.

The table contains detailed information in these areas:

- **Virtual Networks:** Provides information on logical DN, region, primary CIDR, drift type, description of the issue, and the recommended action to resolve it.
- **CIDRs:** Provides information on logical DN, region, CIDR block range, whether it is a primary CIDR or not, the subnets within the CIDR, drift type, description of the issue, and the recommended action to resolve it.
- **Subnets:** Provides information on logical DN, region, IP address, drift type, description of the issue, and the recommended action to resolve it.

For each configuration drift, the **Status** field will indicate the severity and recommended action:

- **Transient (low):** Drifts that are likely due to recent configuration changes. We recommend waiting for the fabric to stabilize and the drift will likely resolve on its own after the next configuration refresh.
  - **Presumed (medium):** Drifts that may or may not be transient. We recommend monitoring the status and troubleshoot the configuration should the drift persist.
  - **Raised (high):** Critical drifts. We recommend verifying the configuration on Cloud APIC and checking for any associated faults. Redeploying the configuration may help resolve communication issues between the Cloud APIC and cloud services. If the issue persists, check the tech-support logs.
- The **Presented Cloud Resources** table shows the information about all the resources that were properly configured in your cloud, split into the same hierarchies that is shown in the **Configuration Drifts** table (Virtual Networks, CIDRs, and Subnets). This table is designed to provide you with better visibility into what virtual networks, CIDRs, and subnets are associated with a specific VRF in your cloud.

## Configuration Drift Troubleshooting

This section provides a few useful command to verify that the configuration drift processes are up and running on your Cloud APIC, check the application logs, and if necessary generate tech support information.

**Step 1** Log in to the Cisco Cloud APIC via console as a `root` user.

**Step 2** Check the status of the configuration drift application.

```
ACI-Cloud-Fabric-1# moquery -d pluginContr/plugin-Cisco_CApicDrift | egrep "dn |pluginSt |operSt |version"
dn: pluginContr/plugin-Cisco_CApicDrift
```

```
operSt: active
pluginSt: active
Verison: 5.1.0
```

**Step 3** Check the status of the application container.

```
ACI-Cloud-Fabric-1# docker ps | grep drift
CONTAINER ID IMAGE COMMAND CREATED STATUS
NAMES
649af6feb72c a5ea08bbf541 "/opt/bin/conit.bi..." 13 hours ago Up 13
hours drift-api-b703e569-0aa6-859f-c538-a5fecbc5708f
```

**Step 4** Check memory consumed by all Docker containers.

Total amount of memory consumed must be under 12GB.

```
ACI-Cloud-Fabric-1# systemctl status ifc-scheduler_allocations.slice | grep Memory
```

**Step 5** If necessary, collect the tech support logs.

Logs will be saved in the /data/techsupport directory on the controller.

```
ACI-Cloud-Fabric-1# trigger techsupport controllers application CApicDrift
ACI-Cloud-Fabric-1# trigger techsupport controllers application CApicDrift vendorName Cisco
```

**Step 6** Check the application logs.

The logs for configuration drift process are stored in the /data2/logs/Cisco\_CApicDrift directory.

The runhist.log file provides information about each time the application was started, for example:

```
cat runhist.log
1 - Thu Jun 11 23:55:59 UTC 2020
2 - Fri Jun 12 01:19:41 UTC 2020
```

The drift.log file is the application log file and can be used to view the number of times configuration drift was updated and how long each update took.

```
cat drift.log | grep ITER
{"file":"online_snapshot.go:178","func":"Wait","level":"info","msg":"ITER# 109
ENDED === RDFGEN TIME: 1m40.383751649s, MODEL UPLOAD TIME 5m54.245550374s; TOTAL
TIME:: 7m34.629447083s","time":"2020-06-12T19:53:13Z"}
```



## CHAPTER 10

# AWS Transit Gateway on Cisco Cloud APIC

---

- [AWS Transit Gateway on Cisco Cloud APIC, on page 185](#)

## AWS Transit Gateway on Cisco Cloud APIC

Beginning in Cisco Cloud Application Policy Infrastructure Controller (APIC) Release 5.0(1), you can use Amazon Web Services (AWS) Transit Gateway with Cisco Cloud APIC. AWS Transit Gateway is a service that functions as an internal router to automate connectivity between virtual private clouds (VPCs). The VPCs can be in different AWS regions in a cloud site.

Virtual private clouds (VPC) can't communicate with each other without additional configuration. Without using AWS Transit Gateway, you can configure inter-VPC communication by configuring VPC peering. Alternatively, you can use VPN tunnels and CCRs.

However, when you use AWS Transit Gateway with Cisco Cloud APIC, you connect VPCs or VRFs in the cloud site simply by associating the VPCs or VRFs to the same AWS Transit Gateways.

Using AWS Transit Gateway with Cisco Cloud APIC provides several benefits: higher performance, simplicity, scalability and potential lower cost.



---

**Note** You can attach a Cisco Cloud APIC user tenant's VPC (CtxProfile) to an AWS Transit Gateway (hub network) only if you have administrator privileges and the user is part of security domain "all". Without such access, you cannot attach the user tenant's VPC to an AWS Transit Gateway.

---

For detailed information about using AWS Transit Gateway with Cisco Cloud APIC, see [Increasing Bandwidth Between VPCs by Using AWS Transit Gateway](#).





# APPENDIX **A**

## Cisco Cloud APIC Error Codes

- [Cisco Cloud APIC Error Codes, on page 187](#)

### Cisco Cloud APIC Error Codes

This section describes the Cisco Cloud APIC error codes.

**Table 39: Cisco Cloud APIC Error Codes**

| Component      | Error Code                                  | Constraint                                                                                                 |
|----------------|---------------------------------------------|------------------------------------------------------------------------------------------------------------|
| cloud-template | CT_INFRANETWORK_COUNT                       | The count of the cloudtemplateInfraNetwork MO is at most 1                                                 |
| cloud-template | CT_INFRANETWORK_COUNT                       | The count of the cloudtemplateInfraNetwork MO is at most 1                                                 |
| cloud-template | CT_INFRANETWORK_VRF                         | In the cloudtemplateInfraNetwork MO, the vrfName must be overlay-1                                         |
| cloud-template | CT_INFRANETWORK_PARENT                      | For the cloudtemplateInfraNetwork MO, the parent MO must be uni/tn-infra                                   |
| cloud-template | CT_INFRANETWORK_NUMROUTERSPERREGION_MINIMUM | In the cloudtemplateInfraNetwork MO, for the attribute numRoutersPerRegion, the minimum allowed value is 2 |

| Component      | Error Code                                      | Constraint                                                                                                                                                                                                    |
|----------------|-------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cloud-template | CT_INFRANETWORK_NUMROUTERSPERREGION_MAXIMUM     | In the <code>cloudtemplateInfraNetwork</code> MO, for the attribute <code>numRoutersPerRegion</code> , the maximum allowed value is 4                                                                         |
| cloud-template | CT_INFRANETWORK_NUMREMOTESITESUBNETPOOL_MINIMUM | In the <code>cloudtemplateInfraNetwork</code> MO, for the attribute <code>numRemoteSiteSubnetPool</code> , the minimum allowed value is 2                                                                     |
| cloud-template | CT_INFRANETWORK_NUMREMOTESITESUBNETPOOL_MAXIMUM | In the <code>cloudtemplateInfraNetwork</code> MO, for the attribute <code>numRemoteSiteSubnetPool</code> , the maximum allowed value is 2                                                                     |
| cloud-template | CT_INTNETWORK_COUNT                             | The count of the <code>cloudtemplateIntNetwork</code> MO is at most 1                                                                                                                                         |
| cloud-template | CT_EXTNETWORK_COUNT                             | The count of the <code>cloudtemplateExtNetwork</code> MO is at most 1                                                                                                                                         |
| cloud-template | CT_VPNNETWORK_COUNT                             | The count of the <code>cloudtemplateVpnNetwork</code> MO is at most 1                                                                                                                                         |
| cloud-template | CT_OSPF_COUNT                                   | The count of the <code>cloudtemplateOspf</code> MO is at most 1                                                                                                                                               |
| cloud-template | CT_INTNETWORK_REGION_MATCH                      | The regions specified by <code>cloudRegionName</code> under <code>cloudtemplateIntNetwork</code> must have a corresponding <code>cloudRegion</code> under <code>cloudProvP</code>                             |
| cloud-template | CT_INTNETWORK_REGION_MANAGED                    | The regions specified by the <code>cloudRegionName</code> children of <code>cloudtemplateIntNetwork</code> must have corresponding <code>cloudRegion</code> with <code>adminSt</code> as <code>managed</code> |

| Component      | Error Code                               | Constraint                                                                                                                                                                                                                                |
|----------------|------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cloud-template | CT_INTNETWORK_REGION_MAXIMUM             | The maximum number of regions ( <code>cloudRegionName</code> ) specified under <code>cloudtemplateIntNetwork</code> is 4                                                                                                                  |
| cloud-template | CT_EXTNETWORK_REGION_SUBSET              | The regions that are specified by the <code>cloudRegionName</code> children of <code>cloudtemplateExtNetwork</code> must also be specified by <code>cloudRegionName</code> children under <code>cloudtemplateIntNetwork</code>            |
| cloud-template | CT_EXTSUBNETPOOL_COUNT                   | The count of the <code>cloudtemplateExtSubnetPool</code> is at most 1                                                                                                                                                                     |
| cloud-template | CT_EXTSUBNETPOOL_SUBNETPOOL_ADDRESS      | In <code>cloudtemplateExtSubnetPool</code> , the subnetpool must contain a network address                                                                                                                                                |
| cloud-template | CT_EXTSUBNETPOOL_SUBNETPOOL_IP_VERSION   | In <code>cloudtemplateExtSubnetPool</code> , the subnetpool must contain a IPv4 address                                                                                                                                                   |
| cloud-template | CT_EXTSUBNETPOOL_SUBNETPOOL_ADDRESS_TYPE | In <code>cloudtemplateExtSubnetPool</code> , the subnetpool IP address must not from multicast or loopback address space                                                                                                                  |
| cloud-template | CT_EXTSUBNETPOOL_SUBNETPOOL_MINIMUM_SIZE | In <code>cloudtemplateExtSubnetPool</code> , the subnetpool must be at least /22 (the netmask must be 22 or less)                                                                                                                         |
| cloud-template | CT_INTNETWORK_MISSING_HOME               | If there are any <code>cloudRegionName</code> under <code>cloudtemplateIntNetwork</code> , then one of the <code>cloudRegonName</code> must be associated to a region that is the home region of the cAPIC ( <code>capicDeployed</code> ) |

| Component      | Error Code                           | Constraint                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cloud-template | CT_CLOUD_APICSUBNETPOOL_INSUFFICIENT | There must be enough <code>cloudApicSubnetPool</code> MOs to generate <code>cloudApicSubnet</code> MOs so that all the <code>cloudRegionName</code> MOs specified under <code>cloudtemplateIntNetwork</code> can be associated to a unique <code>cloudApicSubnet</code> MO. The subnets from the <code>cloudApicSubnet</code> MOs are used as the CIDRs in the <code>cloudCtxProfile</code> of the corresponding region. |
| cloud-template | CT_IPSECTUNNEL_PEERADDR_IP_VERSION   | In <code>cloudtemplateIpSecTunnel</code> , the <code>peeraddr</code> must contain a IPv4 address                                                                                                                                                                                                                                                                                                                         |
| cloud-template | CT_IPSECTUNNEL_PEERADDR_IS_HOST      | In <code>cloudtemplateIpSecTunnel</code> , the <code>peeraddr</code> must be host address (i.e. /32)                                                                                                                                                                                                                                                                                                                     |
| cloud-template | CT_PROFILE_COUNT                     | The count of the <code>cloudtemplateProfile</code> MO is at most 1                                                                                                                                                                                                                                                                                                                                                       |
| cloud-template | CT_PROFILE_DELETE                    | The <code>cloudtemplateProfile</code> MO cannot be deleted unless its parent <code>cloudtemplateInfraNetwork</code> is also deleted                                                                                                                                                                                                                                                                                      |
| cloud-template | CT_PROFILE_ROUTERUSERNAME_NONEMPTY   | In <code>cloudtemplateProfile</code> , the <code>routerUsername</code> must be non-empty                                                                                                                                                                                                                                                                                                                                 |
| cloud-template | CT_PROFILE_ROUTERPASSWORD_NONEMPTY   | In <code>cloudtemplateProfile</code> , the <code>routerPassword</code> must be non-empty                                                                                                                                                                                                                                                                                                                                 |



| Component      | Error Code                         | Constraint                                                                                                                                                                                                                                                                                                   |
|----------------|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cloud-template | CT_PROFILE_ROUTERUSERNAME_MODIFY   | In <code>cloudtemplateProfile</code> , the <code>routerUsername</code> cannot be modified when there are routers deployed in any region, i.e. any <code>cloudRegionName</code> under <code>cloudtemplateIntNetwork</code> (The modification is allowed when there are no router deployments in any region)   |
| cloud-template | CT_PROFILE_ROUTERPASSWORD_MODIFY   | In <code>cloudtemplateProfile</code> , the <code>routerPassword</code> cannot be modified when there are routers deployed in any region, i.e. any <code>cloudRegionName</code> under <code>cloudtemplateIntNetwork</code> (The modification is allowed when there are no router deployments in any region)   |
| cloud-template | CT_PROFILE_ROUTERTHROUGHPUT_MODIFY | In <code>cloudtemplateProfile</code> , the <code>routerThroughput</code> cannot be modified when there are routers deployed in any region, i.e. any <code>cloudRegionName</code> under <code>cloudtemplateIntNetwork</code> (The modification is allowed when there are no router deployments in any region) |
| cloud          | CT_APICSUBNET_INVALID_HOME_REGION  | In a <code>cloudApicSubnet</code> MO, the region marked for <code>capicDeployed</code> must be a valid region                                                                                                                                                                                                |
| cloud          | CT_APICSUBNET_REPEATED_REGION      | In a <code>cloudApicSubnet</code> MO, a region can be associated with at most 1 subnet                                                                                                                                                                                                                       |

| Component | Error Code                             | Constraint                                                                                                                           |
|-----------|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| cloud     | CT_APICSUBNET_MULTIPLE_HOME_REGION     | In <code>cloudApicSubnet</code> MOs, at most, 1 region may have <code>capicDeployed</code> as true                                   |
| cloud     | CLOUD_APICSUBNETPOOL_CREATEDBY_USER    | In <code>cloudApicSubnetPool</code> , the <code>createdBy</code> attribute must be USER                                              |
| cloud     | CLOUD_APICSUBNETPOOL_SUBNET_IP_VERSION | In <code>cloudApicSubnetPool</code> , the subnet must contain a IPv4 address                                                         |
| cloud     | CLOUD_APICSUBNETPOOL_SUBNET_SIZE       | In <code>cloudApicSubnetPool</code> , the subnet must be /24                                                                         |
| cloud     | CLOUD_APICSUBNETPOOL_DELETE_USAGE      | A <code>cloudApicSubnetPool</code> cannot be deleted if at least one of its <code>cloudApicSubnet</code> child is in use by a region |
| cloud     | CLOUD_APICSUBNETPOOL_DELETE_CREATEDBY  | A <code>cloudApicSubnetPool</code> whose <code>createdBy</code> attribute is not USER cannot be deleted                              |