



## **Cisco Cloud APIC for Azure User Guide, Release 5.0(x)**

**First Published:** 2020-05-14

**Last Modified:** 2020-07-03

### **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

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)

© 2020–2022 Cisco Systems, Inc. All rights reserved.



## CONTENTS

---

<b>CHAPTER 1</b>	<b>New and Changed Information 1</b>
	New and Changed Information 1

---

<b>CHAPTER 2</b>	<b>About Cisco Cloud APIC 3</b>
	Overview 3
	Guidelines and Limitations 4
	About the Cisco Cloud APIC GUI 6
	Understanding the Cisco Cloud APIC GUI Icons 6

---

<b>CHAPTER 3</b>	<b>Cisco Cloud APIC Policy Model 13</b>
	About the ACI Policy Model 13
	Policy Model Key Characteristics 13
	Logical Constructs 14
	The Cisco ACI Policy Management Information Model 15
	Tenants 16
	VRFs 17
	Cloud Application Profiles 18
	Cloud Endpoint Groups 19
	Contracts 21
	Comma-separated Filters Support for Contract Rule Consolidation 22
	Filters and Subjects Govern Cloud EPG Communications 23
	About the Cloud Template 24
	Managed Object Relations and Policy Resolution 27
	Default Policies 28
	Shared Services 29

---

<b>CHAPTER 4</b>	<b>Configuring Cisco Cloud APIC Components</b>	<b>31</b>
	About Configuring the Cisco Cloud APIC	31
	Configuring the Cisco Cloud APIC Using the GUI	31
	Creating a Tenant Using the Cisco Cloud APIC GUI	31
	Creating an Application Profile Using the Cisco Cloud APIC GUI	35
	Creating a VRF Using the Cisco Cloud APIC GUI	35
	Creating an EPG Using the Cisco Cloud APIC GUI	36
	Creating a Filter Using the Cisco Cloud APIC GUI	41
	Creating a Contract Using the Cisco Cloud APIC GUI	43
	Creating an Inter-Tenant Contract Using the Cisco Cloud APIC GUI	44
	Specifying Consumer and Provider EPGs Using the Cisco Cloud APIC	47
	Creating a Cloud Context Profile Using the Cisco Cloud APIC GUI	48
	Configuring Virtual Machines in Azure	50
	Creating a Backup Configuration Using the Cisco Cloud APIC GUI	51
	Creating a Tech Support Policy Using the Cisco Cloud APIC GUI	54
	Creating a Scheduler Using the Cisco Cloud APIC GUI	55
	Creating a Remote Location Using the Cisco Cloud APIC GUI	57
	Creating a Login Domain Using the Cisco Cloud APIC GUI	58
	Creating a Security Domain Using the Cisco Cloud APIC GUI	60
	Creating a Role Using the Cisco Cloud APIC GUI	60
	Creating an RBAC Rule Using the Cisco Cloud APIC GUI	65
	Creating a Certificate Authority Using the Cisco Cloud APIC GUI	66
	Creating a Key Ring Using the Cisco Cloud APIC GUI	67
	Creating a Local User Using the Cisco Cloud APIC GUI	69
	Managing Regions (Configuring a Cloud Template) Using the Cisco Cloud APIC GUI	71
	Configuring Smart Licensing	72
	Cloud Resources Naming	73
	Variables Available for Naming Rules	74
	Naming Rules Guidelines and Limitations	75
	Viewing Cloud Resource Naming Rules	76
	Configuring Cisco Cloud APIC Using the REST API	77
	Creating a Tenant Using the REST API	77
	Creating a Contract Using the REST API	78



Creating a Cloud Context Profile Using the REST API	78
Managing a Cloud Region Using the REST API	79
Creating a Filter Using the REST API	79
Creating an Application Profile Using the REST API	80
Creating a Cloud EPG Using the REST API	80
Creating an External Cloud EPG Using the REST API	81
Creating a Cloud Template Using the REST API	82
Defining Global Cloud Resource Naming Rules or Overriding Specific Object's Name	83

**CHAPTER 5****Viewing System Details 85**

Viewing Application Management Details	85
Viewing Cloud Resource Details	86
Viewing Operations Details	88
Viewing Infrastructure Details	90
Viewing Administrative Details	90
Viewing Health Details Using the Cisco Cloud APIC GUI	92

**CHAPTER 6****Deploying Layer 4 to Layer 7 Services 95**

Overview	95
About Service Graphs	95
About Application Load Balancers	96
About Network Load Balancer	97
Dynamic Server Attachment to Server Pool	98
About Inter-VNet Services	98
About Multinodes	99
About Layer 4 to Layer 7 Service Redirect	99
Passthrough Rules	101
Redirect Programming	101
Redirect Policy	102
Workflow for Configuring Redirect	102
Example Use Cases	103
Guidelines and Limitations for Redirect	116
Adding a New CIDR to Overlay-2 Using the Cloud APIC GUI	118
Deploying a Service Graph	119

Deploying a Service Graph Using the GUI	120
Creating Service Devices Using The Cloud APIC GUI	120
Creating a Service Graph Template Using the Cisco Cloud APIC GUI	123
Deploying Layer 4 to Layer 7 Services Using the Cisco Cloud APIC GUI	126
Deploying a Service Graph Using the REST API	130
Creating an Internal-Facing Load Balancer Using the REST API	130
Configuring an Internet-Facing Load Balancer Using the REST API	131
Creating a Third-Party Firewall Using the REST API	132
Creating a Service Graph Using the REST API	133
Creating a Multi-Node Service Graph Using the REST API	134
Creating a Multi-Node Service Graph With Redirect Using the REST API	137
Attaching a Service Graph Using the REST API	141
Configuring an HTTP Service Policy Using the REST API	142
Configuring a Key Ring Using the REST API	143
Creating an HTTPS Service Policy Using the REST API	145

---

**CHAPTER 7**
**Cisco Cloud APIC Security 147**

Access, Authentication, and Accounting	147
Configuration	147
Configuring TACACS+, RADIUS, LDAP and SAML Access	148
Overview	148
Configuring Cloud APIC for TACACS+ Access	148
Configuring Cloud APIC for RADIUS Access	149
Configuring a Cisco Secure Access Control Server for RADIUS and TACACS+ Access to the Cloud APIC	151
Configuring LDAP Access	151
Configuring Windows Server 2008 LDAP for APIC Access with Cisco AVPair	151
Configuring Cloud APIC for LDAP Access	151
Configuring Cloud APIC for SAML Access	153
About SAML	153
Configuring Cloud APIC for SAML Access	154
Setting Up a SAML Application in Okta	155
Setting Up a Relying Party Trust in AD FS	155
Configuring HTTPS Access	155

About HTTPS Access	156
Guidelines for Configuring Custom Certificates	156
Configuring a Custom Certificate for Cisco ACI HTTPS Access Using the GUI	156

---

<b>CHAPTER 8</b>	<b>Configuration Drift</b>	<b>159</b>
	Configuration Drift Notifications and Faults	159
	Enabling Configuration Drift Detection	160
	Checking for Missing Contracts Configuration	161
	Configuration Drift Troubleshooting	164

---

<b>APPENDIX A</b>	<b>Cisco Cloud APIC Error Codes</b>	<b>165</b>
	Cisco Cloud APIC Error Codes	165





# CHAPTER 1

## New and Changed Information

This chapter contains the following sections:

- [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 APIC for Cisco APIC Release 5.0(2)**

Feature or Change	Description	Where Documented
Support for Network (Azure) Load Balancer.	You can deploy a Layer 4 device that distributes the in-bound flow packets to the back-end pool targets.	<a href="#">Deploying Layer 4 to Layer 7 Services, on page 95</a>
Support for Inter-VNET service graphs.	You can deploy a Provider EPG and Service devices in same VPC/VNET.	<a href="#">Deploying Layer 4 to Layer 7 Services, on page 95</a>
Support for Multi-node service Graph	You can enable multiple deployment scenarios with service graphs.	<a href="#">Deploying Layer 4 to Layer 7 Services, on page 95</a>
Support for multiple CIDR and subnet blocks on the infra VNet.	You can configure multiple CIDR and subnet blocks on the infra VNet.	<a href="#">Deploying Layer 4 to Layer 7 Services, on page 95</a>
Support for cloud EPGs and cloud external EPGs in the infra tenant.	You can create cloud EPGs and cloud external EPGs in the infra tenant, where all the cloud EPGs and cloud external EPGs will be associated with the overlay-2 VRF in the infra tenant.	<a href="#">Deploying Layer 4 to Layer 7 Services, on page 95</a>

Feature or Change	Description	Where Documented
Support for Layer 4 to Layer 7 service redirect	Support is now available for Layer 4 to Layer 7 service redirect.	<a href="#">Deploying Layer 4 to Layer 7 Services, on page 95</a>
Tag-based search	While viewing the cloud resource details for Endpoints, search based on cloud tag attribute is supported.	<a href="#">Viewing System Details</a>
Naming convention support for more than 32 characters for the tenant and VRF name combination	All VRFs are assigned a VrfEncoded 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 VrfEncoded value.	<a href="#">Configuring Cisco Cloud APIC Components</a>
Support for comma-separated filters for rule creation in contracts	After a contract is created, some of the rules defined in the contract can be consolidated based on certain criteria.	<a href="#">Cisco Cloud APIC Policy Model</a>
Support for a static IP address for a load balancer	While creating a device, you can assign a static IP address for an application load balancer (ALB) or a network load balancer (NLB).	<a href="#">Deploying Layer 4 to Layer 7 Services, on page 95</a>
Configuration drifts information	Cloud APIC provides visibility into any security policy (contract) configuration discrepancy between what you deploy from the Cloud APIC and what is actually configured in the cloud site.	<a href="#">Configuration Drift Notifications and Faults, on page 159</a>
Custom naming rules for cloud resources	You can create a global naming policy on the Cloud APIC, which allows you to define a custom cloud resources naming convention for all objects deployed from the Cloud APIC into the Azure cloud.	<a href="#">Cloud Resources Naming, on page 73</a>

**Table 2: New Features and Changed Behavior in Cisco APIC for Cisco APIC Release 5.0(1)**

Feature or Change	Description	Where Documented
Additional error codes	Additional error codes have been added as part of Release 5.0(1)	<a href="#">Cisco Cloud APIC Error Codes, on page 165</a>



## CHAPTER 2

# About Cisco Cloud APIC

---

- [Overview, on page 3](#)
- [Guidelines and Limitations, on page 4](#)
- [About the Cisco Cloud APIC GUI, on page 6](#)

## 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). Release 4.1(1) supports Amazon Web Services. Beginning in Release 4.2(x), support is added for Azure.

When deployed, the Cisco Cloud APIC:

- Provides an interface that is similar to the existing Cisco APIC to interact with the Azure 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 Multi-Site Orchestrator to the on-premises and cloud sites, and Cisco Cloud APIC translates the policies to the cloud native constructs 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.

## Guidelines and Limitations

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

- You cannot stretch more than one VRF between on-prem and the cloud while using inter-VRF route leaking in the cloud CSRs (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.
- Before configuring an object for a tenant, first check for any stale cloud resource objects. A stale configuration might be present if it was not cleaned properly from the previous Cisco Cloud APIC virtual machines that managed the account. Cisco Cloud APIC can display stale cloud objects, but it cannot remove them. You must log in to the cloud account and remove them manually.




---

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

Azure allows multiple tenants to share an Azure account owned by one tenant. When the account is shared by multiple tenants, only the owner tenant is able to view the stale objects in the other tenants.

---

To check for stale cloud resources:

1. From the Cisco Cloud APIC GUI, 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.
- Cisco Cloud APIC tries to manage the Azure 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 Azure IAM users in the Azure infra tenant subscription, and the other tenant subscriptions, do not disturb the resources that Cisco Cloud APIC creates. For this purpose, all resources Cisco Cloud APIC creates on Azure has at least one of these two tags:



- AciDnTag
- AciOwnerTag

Cisco Cloud APIC must prevent Azure IAM users who have access to create, delete, or update VM, 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 subscriptions. Azure subscription 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:

```
{
  "properties": {
    "level": "CanNotDelete",
    "notes": "Optional text notes."
  }
}
```

- 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.
- 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.

**Table 3. Scale Supported**

Component	Number Supported
Tenants	20
Application Profiles	500

Component	Number Supported
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**, **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 EPG-specific window, hover your mouse over the **Navigation** menu and click **Application Management > EPGs**. From there, you can use the **Navigation** menu to view the details of another component. For example, you can navigate to the **Active Sessions** window from **EPGs** by clicking **Operations > Active Sessions**.

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 **Routers** 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.

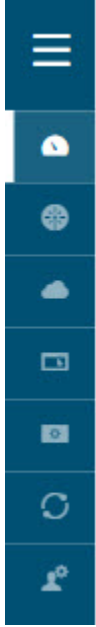
For more information about the GUI icons, see [Understanding the Cisco Cloud APIC GUI Icons, on page 6](#)

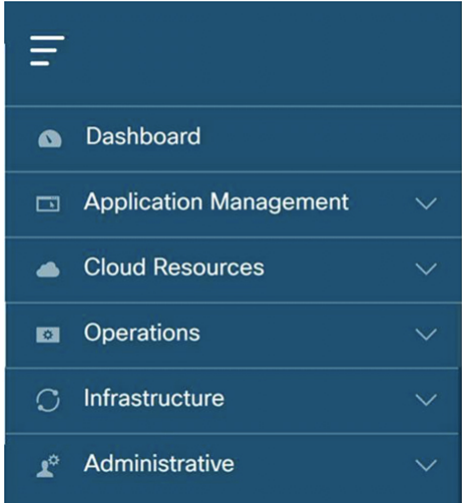

For more information about configuring Cisco Cloud APIC components, see [Configuring Cisco Cloud APIC Components, on page 31](#)

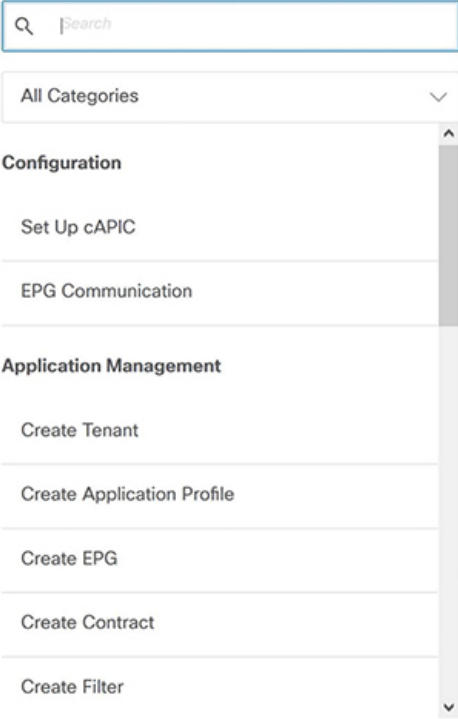
## Understanding the Cisco Cloud APIC GUI Icons

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





Table 4: Cisco Cloud APIC GUI Icons

Icon	Description
<p data-bbox="386 344 716 371"><i>Figure 1: Navigation Pane (Collapsed)</i></p> 	<p data-bbox="938 344 1523 630">The left side of the GUI contains the <b>Navigation</b> 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 <b>Navigation</b> pane locks in the open position. To collapse it, click the menu icon again. When you expand the <b>Navigation</b> pane by hovering the mouse icon over the menu icon, you collapse the <b>Navigation</b> pane by moving the mouse icon away from it.</p> <p data-bbox="938 646 1523 772">When expanded, the <b>Navigation</b> 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="349 289 675 315"><b>Figure 2: Navigation Pane (Expanded)</b></p> 	<p data-bbox="901 289 1417 352">The Cisco Cloud APIC component windows are organized in the <b>Navigation</b> pane as follows:</p> <ul data-bbox="938 373 1482 1136" style="list-style-type: none"> <li>• <b>Dashboard</b> Tab—Displays summary information about the Cisco Cloud APIC components.</li> <li>• <b>Application Management</b> Tab—Displays information about tenants, application profiles, EPGs, contracts, filters, VRFs, service graphs, devices, and cloud context profiles.</li> <li>• <b>Cloud Resources</b> Tab—Displays information about regions, VNETs, routers, security groups (application security groups/network security groups), endpoints, instances, and cloud services (and target groups).</li> <li>• <b>Operations</b> Tab—Displays information about event analytics, active sessions, backup &amp; restore policies, tech support policies, firmware management, schedulers, and remote locations.</li> <li>• <b>Infrastructure</b> Tab—Displays information about the system configuration, inter-region connectivity, and on-premises connectivity.</li> <li>• <b>Administrative</b> Tab—Displays information about authentication, event analytics, security, local and remote users, and smart licensing.</li> </ul> <p data-bbox="901 1171 1482 1262"><b>Note</b> For more information about the contents of these tabs, see <a href="#">Viewing System Details, on page 85</a></p>
<p data-bbox="349 1306 613 1331"><b>Figure 3: Intent Menu-Bar Icon</b></p> 	<p data-bbox="901 1306 1466 1369">The <b>Intent</b> icon appears in the menu bar between the <b>search</b> and the <b>help</b> icons.</p> <p data-bbox="901 1390 1482 1577">When clicked, the <b>Intent</b> dialog appears (see below). The <b>Intent</b> 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 <b>Intent</b> icon. Close the dialog box to access the <b>Intent</b> icon again.</p> <p data-bbox="901 1598 1482 1688">For more information about creating a component, see <a href="#">Configuring Cisco Cloud APIC Components, on page 31</a>.</p>

Icon	Description
<p><b>Figure 4: Intent Dialog Box</b></p> 	

Icon	Description
	<p>The <b>Intent</b> 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> <ul style="list-style-type: none"> <li>• <b>All Categories</b></li> <li>• <b>Configuration</b>—Displays the following options: <ul style="list-style-type: none"> <li>• <b>Set Up cAPIC</b></li> <li>• <b>EPG Communication</b></li> </ul> </li> <li>• <b>Application Management</b>—Displays the following options: <ul style="list-style-type: none"> <li>• <b>Create Tenant</b></li> <li>• <b>Create Application Profile</b></li> <li>• <b>Create EPG</b></li> <li>• <b>Create Contract</b></li> <li>• <b>Create Filter</b></li> <li>• <b>Create VRF</b></li> <li>• <b>Create Device</b></li> <li>• <b>Create Service Graph</b></li> <li>• <b>Create Cloud Context Profile</b></li> </ul> </li> <li>• <b>Operations</b>—Displays the following options: <ul style="list-style-type: none"> <li>• <b>Create Backup Configuration</b></li> <li>• <b>Create Tech Support</b></li> <li>• <b>Create Scheduler</b></li> <li>• <b>Create Remote Location</b></li> </ul> </li> <li>• <b>Administrative</b>—Displays the following options: <ul style="list-style-type: none"> <li>• <b>Create Login Domain</b></li> <li>• <b>Create Security Domain</b></li> <li>• <b>Create Role</b></li> <li>• <b>Create RBAC Rule</b></li> <li>• <b>Create Certificate Authority</b></li> <li>• <b>Create Key Ring</b></li> </ul> </li> </ul>

Icon	Description
	<ul style="list-style-type: none"> <li>• <b>Create Local User</b></li> </ul>
<p data-bbox="386 373 643 399"><b>Figure 5: Help Menu-Bar Icon</b></p> 	<p>The <b>help</b> menu-bar icon opens the <a href="#">Cisco Cloud APIC Quick Start Guide</a> .</p>
<p data-bbox="386 518 716 543"><b>Figure 6: System Tools Menu-Bar Icon</b></p> 	<p>The <b>system tools</b> menu-bar icon provides the following options:</p> <ul style="list-style-type: none"> <li>• <b>About</b>—Display the Cisco Cloud APIC version.</li> <li>• <b>ObjectStore Browser</b>—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.</li> </ul>
<p data-bbox="386 850 667 875"><b>Figure 7: Search Menu-Bar Icon</b></p> 	<p>The <b>search</b> 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="386 1050 708 1075"><b>Figure 8: User Profile Menu-Bar Icon</b></p> 	<p>The <b>user profile</b> menu-bar icon provides the following options:</p> <ul style="list-style-type: none"> <li>• <b>Change Password</b>—Enables you to change the password.</li> <li>• <b>Change SSH Key</b>—Enables you to change the SSH key.</li> <li>• <b>Change User Certificate</b>—Enables you to change the user certificate.</li> <li>• <b>Logout</b>—Enables you to log out of the GUI.</li> </ul>







## CHAPTER 3

# Cisco Cloud APIC Policy Model

---

- [About the ACI Policy Model, on page 13](#)
- [Policy Model Key Characteristics, on page 13](#)
- [Logical Constructs, on page 14](#)
- [The Cisco ACI Policy Management Information Model, on page 15](#)
- [Tenants, on page 16](#)
- [VRFs, on page 17](#)
- [Cloud Application Profiles, on page 18](#)
- [Cloud Endpoint Groups, on page 19](#)
- [Contracts, on page 21](#)
- [About the Cloud Template, on page 24](#)
- [Managed Object Relations and Policy Resolution, on page 27](#)
- [Default Policies, on page 28](#)
- [Shared Services, on page 29](#)

## 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, 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 9: 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 10: 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.

---

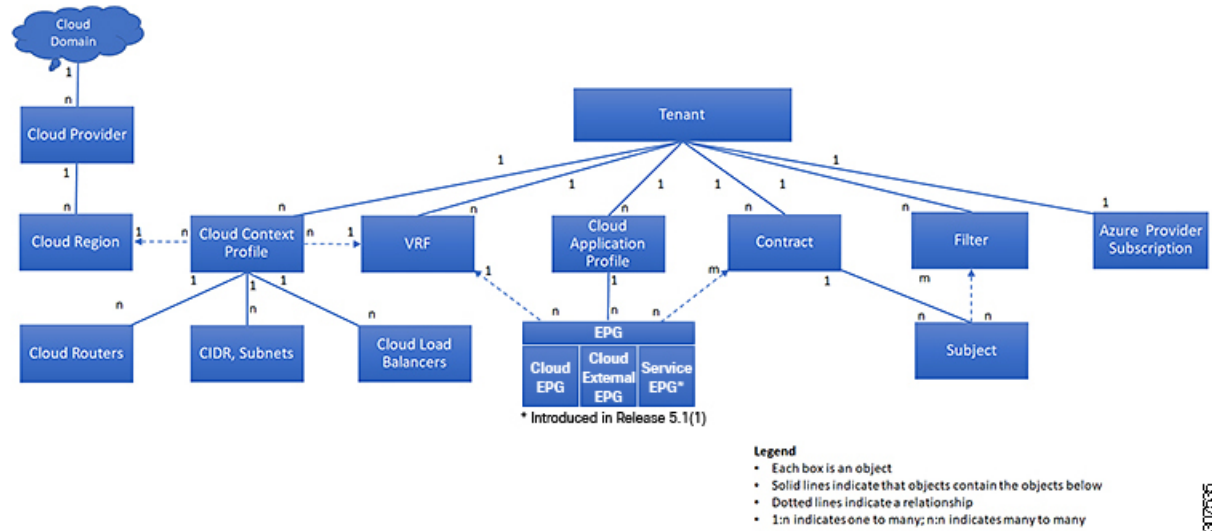
- The infrastructure tenant is provided by the system but can be configured by the cloud infrastructure administrator. It contains policies that govern the operation of infrastructure resources. It also enables a cloud infrastructure provider to selectively deploy resources to one or more user tenants. 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 95](#)
- 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 147](#)

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 11: Tenants



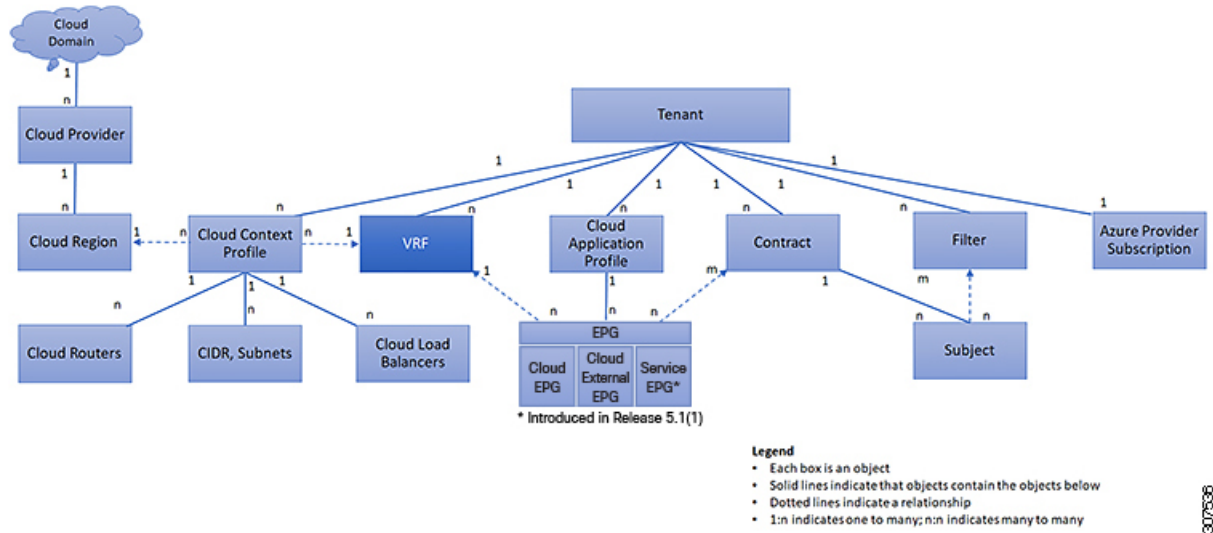
Tenants can be isolated from one another or can share resources. The primary elements that the tenant contains are filters, contracts, Virtual Routing and Forwarding (VRF) instances, cloud context profiles, Azure provider configurations, and cloud application profiles that contain cloud endpoint groups (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, tenant and region, represents a resource group in Azure. A VNET is created inside the resource group based on the VRF name.

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.

## VRFs

A Virtual Routing and Forwarding (VRF) object (`fVContext`) or context is a tenant network (called a VRF 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 12: 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.

Beginning with Release 5.0(2), you can have a hub VNet (a cloudCtxProfile in the infra tenant) that can be carved out into multiple VRFs.

## Cloud Application Profiles

A cloud application profile (`cloudAp`) 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 13: 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:

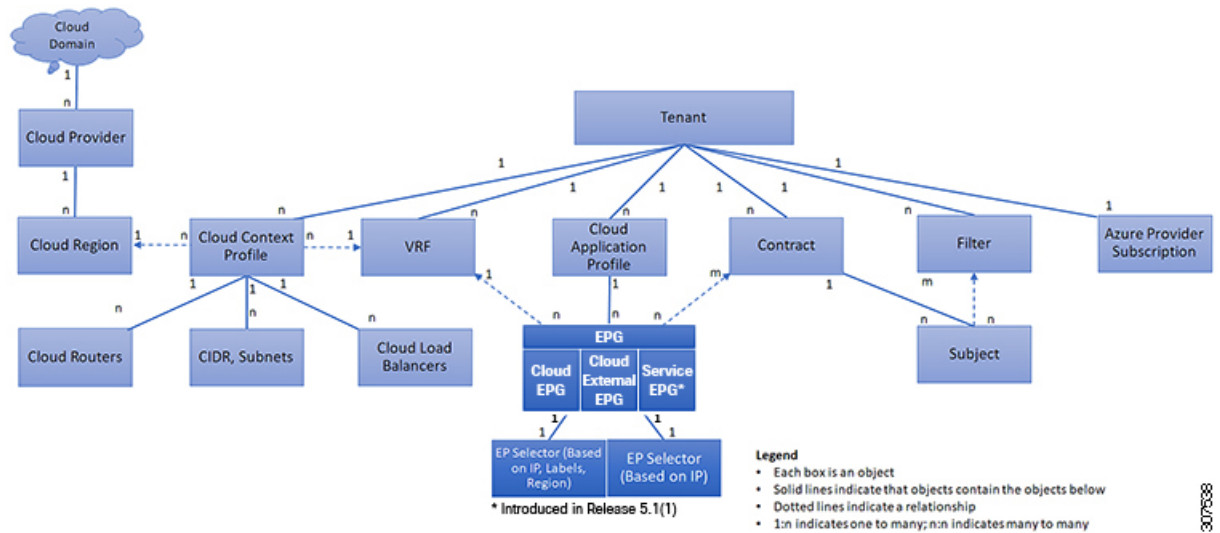
- The application they provide, such as a DNS server or SAP application (see *Tenant Policy Example* in *Cisco APIC REST API Configuration Guide*).
- The function they provide (such as infrastructure)
- Where they are in the structure of the data center (such as DMZ)
- Whatever organizing principle that 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 14: 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. 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 (`cloudEPg`) 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 Azure VNET 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 (vzBrCP) 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 15: 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 the cloud endpoints in that cloud EPG can be initiated from cloud endpoints in 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.

## Comma-separated Filters Support for Contract Rule Consolidation

After a contract is created, some of the rules defined in the contract are consolidated and displayed in Azure based on certain criteria. You can combine multiple ports and multiple IP addresses and ranges into a single, easy-to-understand rule. The criteria for consolidation of rules are:

- Rules are consolidated only within a contract. Two rules resulting from two different contracts are not consolidated in Azure.
- The source/ destination address prefixes and destination port(s) are consolidated.
- The conditions for multiple rules to get consolidated together in an NSG are:
  - Same contract
  - Same protocol (UDP, TCP, ICMP)
  - Same direction (inbound , outbound)
  - Same type (SG, IP)
- Overlapping port ranges for same protocol (TCP/UDP) in the same contract are consolidated to one range.
 

For example, TCP ports 100-200, 150-250 are consolidated to 100-250.
- If 1.2.3.4/32 (any address prefixes) is allowed, and an ext EPG with 0.0.0.0/0 is added, then the allowed Source/Destination IP would be *Any*, not [1.2.3.4/32, 0.0.0.0/0].

Example below shows the EPG1 outbound rules and the consolidated EPG1 outbound rules, based on contracts C1 and C2.

```
Contract C1:
Consumer: EPG1 , Provider: EPG2
Filter: TCP (ports 53)
Filter: UDP (port 53, 5000)
```

```
Contract C2:
Consumer: EPG1 , Provider: EPG2
Filter: TCP (ports 80, 8080)
```

```
EPG1 outbound rules:
EPG1 -> EPG2   TCP   80
EPG1 -> EPG2   TCP  8080
EPG1 -> EPG2   TCP           53
EPG1 -> EPG2   UDP   53
EPG1 -> EPG2   UDP  5000
EPG1 -> 1.1.1.1/32 TCP   80
EPG1 -> 1.1.1.1/32 TCP  8080
EPG1 -> 1.1.1.1/32 TCP   53
EPG1 -> 1.1.1.1/32 UDP   53
EPG1 -> 1.1.1.1/32 UDP  5000
```

```

EPG1 -> 2.2.2.2/32 TCP 80
EPG1 -> 2.2.2.2/32 TCP 8080
EPG1 -> 2.2.2.2/32 TCP 53
EPG1 -> 2.2.2.2/32 UDP 53
EPG1 -> 2.2.2.2/32 UDP 5000

```

Rules are consolidated by comma-separated filters (consolidated based on C1 and C2):

```

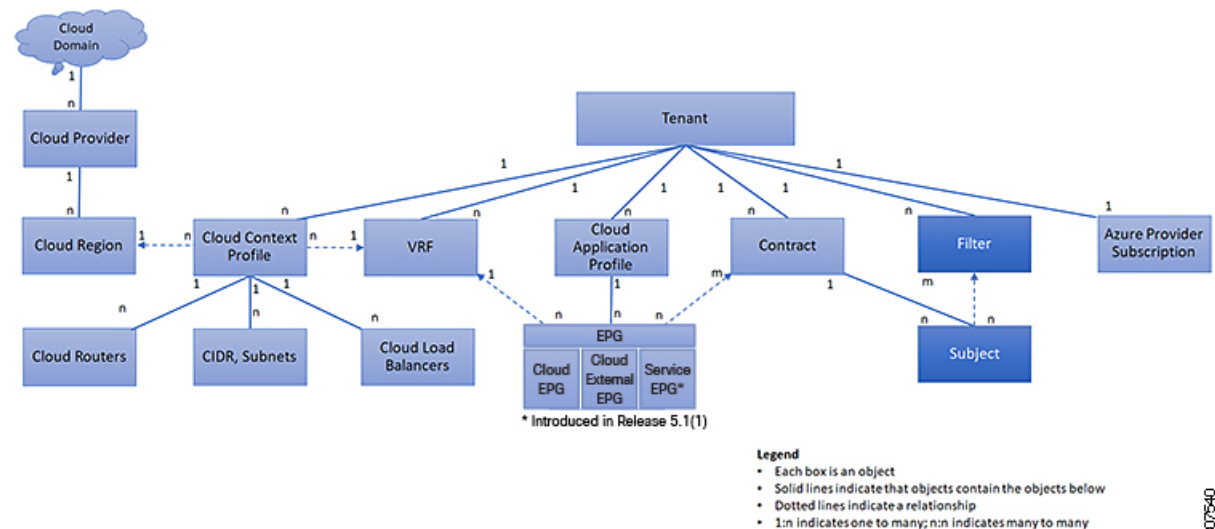
EPG1 -> EPG2 TCP 80,8080
EPG1 -> EPG2 UDP 53,5000
EPG1 -> EPG2 TCP 53
EPG1 -> 1.1.1.1/32, 2.2.2.2/32 TCP 80,8080
EPG1 -> 1.1.1.1/32, 2.2.2.2/32 UDP 53,5000
EPG1 -> 1.1.1.1/32, 2.2.2.2/32 TCP 53

```

## 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 16: 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 Azure, 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 3 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.

- Subjects are contained in contracts. A subject within a contract uses 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.
- ACI contracts rendered in Azure 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 Azure network configuration is the Virtual Private Cloud (VNET). Azure supports many regions worldwide and one VNET is specific to one region.

The cloud template accepts one or more region names and generates the entire configuration for the infra VNETs in those regions. They are the infra VNETs. The Cisco Cloud APIC-managed object (MO) corresponding to the Azure VNET 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. The `cloudCtxProfile` MO for the infra VNet is generated by the cloud template. It carries `ctxProfileOwner == SYSTEM`, which means that this MO is generated by the 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 Azure VNet is the CIDR. In Cisco Cloud APIC, you can choose and deploy CIDRs in the user VNETs. The CIDRs for the infra VNet are provided by users to the cloud template during the initial setup of the cloud site, and are deployed to the Azure cloud by the cloud template.

Beginning with Release 5.0(2), a new property called `createdBy` is added for the CIDR. The default value for this `createdBy` property is `USER`.

- For all user-created CIDRs, the value for the `createdBy` property is set to `USER`.
- For cloud template-created CIDRs, the value for the `createdBy` property is set to `SYSTEM`.

In releases prior to Release 5.0(2), you are not allowed to add more CIDRs to the infra VNet. Beginning with Release 5.0(2), multiple CIDR and subnet blocks can now be configured on the infra VNet. You can create CIDRs and associate subnets in the infra VNet. The cloud template subnets will be mapped to the overlay-1 VRF, but the user-created subnets will be implicitly mapped to the overlay-2 VRF in the same infra VNet. All subnets in the respective VRFs will have separate route tables in the cloud for VRF segregation.

In addition, beginning with Release 5.0(2), you can create cloud EPGs and cloud external EPGs in the infra tenant, where all the cloud EPGs and cloud external EPGs will be associated with the overlay-2 VRF in the infra tenant. A cloud EPG in the overlay-2 VRF can communicate with other cloud EPGs and cloud external

EPGs in the overlay-2 VRF, and can also communicate with cloud EPGs in other user tenant VRFs. We recommend that you do not use existing "cloud-infra" application profiles, and instead create a new application profile in the infra tenant and associate that new application profile to the cloud EPGs and cloud external EPGs in the overlay-2 VRF.

For more information, see [Creating an EPG Using the Cisco Cloud APIC GUI, on page 36](#) and [About the Overlay-1 and Overlay-2 VRFs, on page 100](#).

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

- Subnets
- 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 5: 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"> <li>• <code>routerUsername</code></li> </ul> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>• The username cannot be "admin."</li> <li>• Any username restrictions from Azure applies.</li> </ul> <ul style="list-style-type: none"> <li>• <code>routerPassword</code></li> <li>• <code>routerThroughput</code></li> <li>• <code>routerLicenseToken</code></li> </ul>
<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

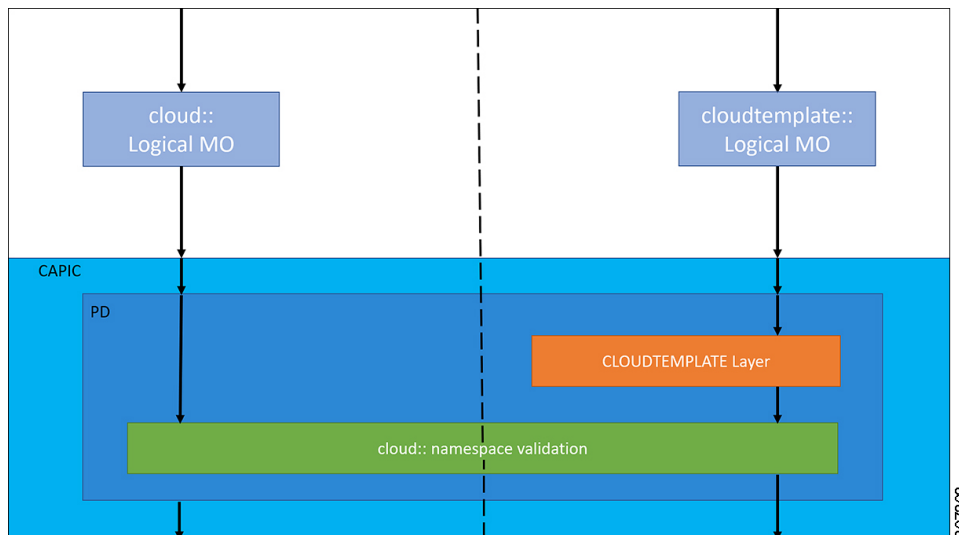
MO	Purpose
cloudtemplateExtNetwork	Contains infra network configuration input that is external of the cloud.  Contains a list of regions where cloud routers are configured for external networking.  Each region is captured through a <code>cloudRegionName</code> child MO
cloudtemplateVpnNetwork	Contains information for setting up a VPN with an ACI on-premises site or another Cisco Cloud APIC site.
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 17: Cloud and Cloud Template MO Conversion**





**Note** For information about configuring the cloud template, see [Configuring Cisco Cloud APIC Components](#), on page 31

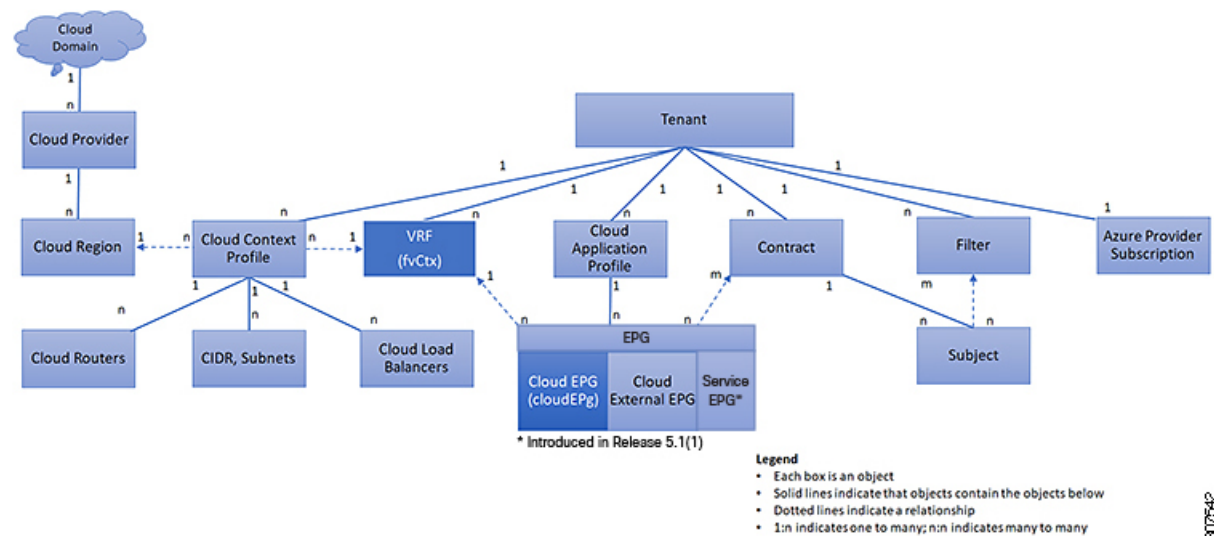
## Managed Object Relations and Policy Resolution

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 `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 18: 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 Azure 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.
  - Inter-tenant contracts require a global scope.





## CHAPTER 4

# Configuring Cisco Cloud APIC Components

- [About Configuring the Cisco Cloud APIC, on page 31](#)
- [Configuring the Cisco Cloud APIC Using the GUI, on page 31](#)
- [Configuring Cisco Cloud APIC Using the REST API, on page 77](#)

## 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 95](#).
- For information about the GUI, such as navigation and a list of configurable components, see [About the Cisco Cloud APIC GUI, on page 6](#).

## Configuring the Cisco Cloud APIC Using the GUI

### Creating a Tenant Using the Cisco Cloud APIC GUI

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

#### Before you begin

- You can create a tenant that is managed by the Cisco Cloud APIC or a tenant that is unmanaged. To establish a managed tenant, you must first obtain the Azure subscription ID from the Azure portal. You enter the subscription ID in the appropriate field of the Cisco Cloud APIC when creating the tenant. Before you can use the managed tenant, you must explicitly grant the Cisco Cloud APIC permission to manage the subscription. The steps for doing so are displayed in the Cisco Cloud APIC GUI during tenant creation. The steps for the infra tenant, however, are displayed in the infra tenant details view:

1. Click the **Navigation** menu > **Application Management** subtab.

2. Double-click the infra tenant.
3. Click **View Azure Role Assignment Command**. The steps for granting the Cisco Cloud APIC permission to manage the subscription are displayed.




---

**Note** For information about obtaining the Azure subscription ID, see the Microsoft Azure documentation.

---

- Creating an unmanaged tenant requires obtaining a directory (Azure Tenant) ID, an Azure enterprise application ID, and a client secret from the enterprise application. For more information, see the Microsoft Azure documentation.




---

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

---

- The required steps to explicitly grant the Cisco Cloud APIC permission to manage a given subscription are located in the Cisco Cloud APIC GUI. When creating a tenant, the steps are displayed after entering the client secret. For the infra tenant:
- 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 Azure subscription 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 Azure subscription 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 Azure Resource Groups. When a new tenant in subscription TA1 in region R2 is managed by Cloud APIC, a Resource Group CAPIC\_TA1\_R2 (e.g. CAPIC\_123456789012\_\_eastus2) is created in the subscription. This Resource Group has a resource tag AciOwnerTag with value IA1\_R1\_TA1\_R2, assuming it was managed by Cloud APIC in subscription 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 a subscription, and then taken down and Cloud APIC is installed in a different subscription. 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 Azure subscription and manually remove the affected Resource Group (for example: CAPIC\_123456789012\_\_eastus2). Next, reload Cloud APIC or delete and add the tenant again.

- 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** Choose the appropriate options and enter the appropriate values in each field as listed in the following *Create Tenant Dialog Box Fields* table then continue.

**Table 6: Create Tenant Dialog Box Fields**

Properties	Description
<b>Name</b>	Enter the name of the tenant.
<b>Description</b>	Enter a description of the tenant.
<b>Settings</b>	
<b>Add Security Domain</b>	To add a security domain for the tenant: <ol style="list-style-type: none"> <li>Click <b>Add Security Domain</b>. The <b>Select Security Domains</b> dialog appears with a list of security domains in the left pane.</li> <li>Click to choose a security domain.</li> <li>Click <b>Select</b> to add the security domain to the tenant.</li> </ol>
<b>Azure Subscription</b>	
<b>Mode</b>	Choose an account type: <ul style="list-style-type: none"> <li>• <b>Create Own</b>—Choose this option to create a new tenant.</li> <li>• <b>Select Shared</b>—Choose this option to inherit the managed or unmanaged settings from an existing tenant.</li> </ul>
<b>Azure Subscription ID</b>	Enter the Azure subscription ID.

Properties	Description
<b>Access Type</b>	Choose an access type: <ul style="list-style-type: none"> <li>• <b>Unmanaged Identity</b>—Choose this option if the tenant subscription is not managed by the Cisco Cloud APIC.</li> <li>• <b>Managed Identity</b>—Choose this option if the tenant subscription is managed by the Cisco Cloud APIC. For more information, see <i>Configuring a Tenant Azure Provider</i>.</li> </ul>
<b>Application ID</b>	<p><b>Note</b> This field is only valid for the <b>Unmanaged Identity</b> access type.</p> <p>Enter the application ID.</p> <p><b>Note</b> For information about obtaining the application ID, see the Azure documentation or support.</p>
<b>Client Secret</b>	<p><b>Note</b> This field is only valid for the <b>Unmanaged Identity</b> access type.</p> <p>Enter the client secret.</p> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>• For information about creating a client secret, see the Azure documentation or support.</li> <li>• You must explicitly grant Cloud APIC permission to manage a given subscription. Go to the Azure portal and follow these steps:               <ol style="list-style-type: none"> <li>a. Open the Cloud Shell</li> <li>b. Choose 'Bash'</li> <li>c. Copy and paste the command displayed in the Cisco Cloud APIC GUI.</li> </ol> </li> </ul>
<b>Active Directory ID</b>	<p><b>Note</b> This field is only valid for the <b>Unmanaged Identity</b> access type.</p> <p>Enter the active directory ID.</p> <p><b>Note</b> For information about obtaining the active directory ID, see the Azure documentation or support.</p>

Properties	Description
Add Security Domain	<p>To add a security domain for the account:</p> <ol style="list-style-type: none"> <li>Click <b>Add Security Domain</b>. The <b>Select Security Domains</b> dialog appears with a list of security domains in the left pane.</li> <li>Click to choose a security domain.</li> <li>Click <b>Select</b> to add the security domain to the tenant.</li> </ol>

**Step 5** Click **Save** when finished.

## 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 7: Create VRF Dialog Box Fields*

Properties	Description
<b>General</b>	
<b>Name</b>	Enter a name for the VRF in the <b>Name</b> 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 <b>Application Management &gt; VRFs</b> subtab. Click a VRF on the right hand pane and look for <i>Encoded VRF Name in Cloud Router</i> .
<b>Tenant</b>	To choose a tenant:  a. Click <b>Select Tenant</b> . The <b>Select Tenant</b> dialog box appears.  b. From the <b>Select Tenant</b> dialog, click to choose a tenant in the left column then click <b>Select</b> . You return to the <b>Create VRF</b> dialog box.
<b>Description</b>	Enter a description of the VRF.

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

## 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.





**Note** Beginning with Release 5.0(2), Cisco Cloud APIC creates the overlay-2 VRF in the infra tenant by default during the bring up, along with the overlay-1 VRF.

In addition, beginning with Release 5.0(2), you can create cloud EPGs and cloud external EPGs in the infra tenant, where all the cloud EPGs and cloud external EPGs will be associated with the overlay-2 VRF in the infra tenant. A cloud EPG in the overlay-2 VRF can communicate with other cloud EPGs and cloud external EPGs in the overlay-2 VRF, and can also communicate with cloud EPGs in other user tenant VRFs. We recommend that you do not use existing "cloud-infra" application profiles, and instead create a new application profile in the infra tenant and associate that new application profile to the cloud EPGs and cloud external EPGs in the overlay-2 VRF.

### 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 8: Create EPG Dialog Box Fields**

Properties	Description
<b>Name</b>	Enter the name of the EPG.
<b>Tenant</b>	<p>To choose a tenant:</p> <ol style="list-style-type: none"> <li>Click <b>Select Tenant</b>. The <b>Select Tenant</b> dialog box appears.</li> <li>From the <b>Select Tenant</b> dialog, click to choose a tenant in the left column. Beginning with Release 5.0(2), you can select the infra tenant and can create cloud EPGs and cloud external EPGs in the infra tenant, as described earlier in this section.</li> <li>Click <b>Select</b>. You return to the <b>Create EPG</b> dialog box.</li> </ol>

Properties	Description
<b>Application Profile</b>	<p>To choose an application profile:</p> <ol style="list-style-type: none"> <li>Click <b>Select Application Profile</b>. The <b>Select Application Profile</b> dialog box appears.</li> <li>From the <b>Select Application Profile</b> dialog, click to choose an application profile in the left column. <ul style="list-style-type: none"> <li><b>Note</b> If you are creating an EPG in the infra tenant, we recommend that you do not choose the <code>cloud-infra</code> application profile because that application profile is used by EPGs in the overlay-1 VRF. Select a different application profile or click <b>Create Application Profile</b> to create a new one.</li> </ul> </li> <li>Click <b>Select</b>. You return to the <b>Create EPG</b> dialog box.</li> </ol>
<b>Description</b>	Enter a description of the EPG.
<b>Settings</b>	
<b>Type</b>	<p>Choose the EPG type:</p> <ul style="list-style-type: none"> <li>• <b>Cloud</b> - Click to create the EPG in the cloud.</li> <li>• <b>External</b> - Click to create an external EPG.</li> </ul>
<b>VRF</b>	<p>To choose a VRF:</p> <ol style="list-style-type: none"> <li>Click <b>Select VRF</b>. The <b>Select VRF</b> dialog box appears.</li> <li>From the <b>Select VRF</b> dialog, click to choose a VRF in the left column. <p>If you are creating an EPG in the infra tenant, select the <b>overlay-2</b> VRF in this step. A cloud EPG in the overlay-2 VRF can communicate with other cloud EPGs and cloud external EPGs in the overlay-2 VRF, and can also communicate with cloud EPGs in other user tenant VRFs.</p> </li> <li>Click <b>Select</b>. You return to the <b>Create EPG</b> dialog box.</li> </ol>

Properties	Description
Endpoint Selectors	

Properties	Description
	<p><b>Note</b> See <a href="#">Configuring Virtual Machines in Azure, on page 50</a> for instructions on configuring virtual machines in Azure as part of the endpoint selector configuration process.</p> <p>To add an endpoint selector:</p> <ol style="list-style-type: none"> <li>a. Click <b>Add Endpoint Selector</b> to open the <b>Add Endpoint Selector</b> dialog.</li> <li>b. In the <b>Add Endpoint Selector</b> dialog, enter a name in the <b>Name</b> field.</li> <li>c. Click <b>Selector Expression</b>. The <b>Key</b>, <b>Operator</b>, and <b>Value</b> fields are enabled.</li> <li>d. Click the <b>Key</b> drop-down list to choose a key. The options are: <ul style="list-style-type: none"> <li>• Choose <b>IP</b> if you want to use an IP address or subnet for the endpoint selector.</li> <li>• Choose <b>Region</b> if you want to use the Azure region for the endpoint selector.</li> <li>• Choose <b>Custom</b> if you want to create a custom key for the endpoint selector.</li> </ul> <p><b>Note</b> When choosing the <b>Custom</b> option, the drop-down list becomes a text box. You need to enter a name for the key in the spaces after <b>custom:</b> (for example, <b>custom: Location</b>).</p> </li> <li>e. Click the <b>Operator</b> drop-down list to choose an operator. The options are: <ul style="list-style-type: none"> <li>• <b>equals</b>: Used when you have a single value in the Value field.</li> <li>• <b>not equals</b>: Used when you have a single value in the Value field.</li> <li>• <b>in</b>: Used when you have multiple comma-separated values in the Value field.</li> <li>• <b>not in</b>: Used when you have multiple comma-separated values in the Value field.</li> <li>• <b>has key</b>: Used if the expression contains only a key.</li> <li>• <b>does not have key</b>: Used if the expression contains only a key.</li> </ul> </li> <li>f. Enter a value in the <b>Value</b> field then click the check mark to validate the entries. The value you enter depends on the choices you made for the <b>Key</b> and <b>Operator</b> fields. For example, if the <b>Key</b> field is set to <b>IP</b> and the <b>Operator</b> field is set to <b>equals</b>, the <b>Value</b> field must be an IP address or subnet. However, if the <b>Operator</b> field is set to <b>has key</b>, the <b>Value</b> field is disabled.</li> <li>g. When finished, click the check mark to validate the selector expression.</li> <li>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.</li> </ol> <p>For example, assume you created two sets of expressions under a single endpoint selector:</p> <ul style="list-style-type: none"> <li>• Endpoint selector 1, expression 1: <ul style="list-style-type: none"> <li>• <b>Key</b>: Region</li> <li>• <b>Operator</b>: equals</li> </ul> </li> </ul>

Properties	Description
	<ul style="list-style-type: none"> <li>• <b>Value:</b> westus</li> </ul> <ul style="list-style-type: none"> <li>• Endpoint selector 1, expression 2:           <ul style="list-style-type: none"> <li>• <b>Key:</b> IP</li> <li>• <b>Operator:</b> equals</li> <li>• <b>Value:</b> 192.0.2.1/24</li> </ul> </li> </ul> <p>In this case, if <i>both</i> of these expressions are true (if the region is westus AND if the IP address belongs to subnet 192.0.2.1/24), then that endpoint is assigned to the Cloud EPG.</p> <p>i. Click the check mark after every additional expression that you want to create under this endpoint selector then click <b>Add</b> 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"> <li>• Endpoint selector 2, expression 1:           <ul style="list-style-type: none"> <li>• <b>Key:</b> Region</li> <li>• <b>Operator:</b> in</li> <li>• <b>Value:</b> eastus, centralus</li> </ul> </li> </ul> <p>In this case:</p> <ul style="list-style-type: none"> <li>• If the region is westus AND the IP address belongs to the 192.0.2.1/24 subnet (endpoint selector 1 expressions)</li> </ul> <p>OR</p> <ul style="list-style-type: none"> <li>• If the region is either eastus or centralus (endpoint selector 2 expression)</li> </ul> <p>Then that end point is assigned to the Cloud EPG.</p>

**Step 5** Click **Save** when finished.

## 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 9: Create Filter Dialog Box Fields**

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

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

## 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 10: Create Contract Dialog Box Fields*

Properties	Description
<b>Name</b>	Enter the name of the contract.
<b>Tenant</b>	<p>To choose a tenant:</p> <ol style="list-style-type: none"> <li>Click <b>Select Tenant</b>. The <b>Select Tenant</b> dialog box appears.</li> <li>From the <b>Select Tenant</b> dialog, click to choose a tenant in the left column. <ul style="list-style-type: none"> <li><b>Note</b> Beginning in Release 5.0(2), you can create contracts in the infra tenant. You can also export contracts from and import contracts to the infra tenant for shared services use cases.</li> </ul> </li> <li>Click <b>Select</b>. You return to the <b>Create Contract</b> dialog box.</li> </ol>
<b>Description</b>	Enter a description of the contract.
<b>Settings</b>	

Properties	Description
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><b>Note</b> 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 <b>Global</b> scope.</p> <p>To enable an EPG in one VRF to communicate with another EPG in a different VRF, choose <b>Global</b> or <b>Tenant</b> scope.</p> <p>For more information about shared services, see <a href="#">Shared Services, on page 29</a>.</p> <p>Click the drop-down arrow to choose from the following scope options:</p> <ul style="list-style-type: none"> <li>• <b>Application Profile</b></li> <li>• <b>VRF</b></li> <li>• <b>Global</b></li> <li>• <b>Tenant</b></li> </ul>
Add Filter	<p>To choose a filter:</p> <ol style="list-style-type: none"> <li>a. Click <b>Add Filter</b>. The filter row appears with a <b>Select Filter</b> option.</li> <li>b. Click <b>Select Filter</b>. The <b>Select Filter</b> dialog box appears.</li> <li>c. From the <b>Select Filter</b> dialog, click to choose a filter in the left column then click <b>Select</b>. You return to the <b>Create Contract</b> dialog box.</li> </ol>

**Step 5** Click **Save** when finished.

## Creating an Inter-Tenant Contract Using the Cisco Cloud APIC GUI

This section explains how to create an inter-tenant contract using the Cisco Cloud APIC GUI. See [Shared Services, on page 29](#) for more information on situations where you might want to create an inter-tenant contract.

### 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 11: Create Contract Dialog Box Fields**

Properties	Description
<b>Name</b>	Enter the name of the contract.
<b>Tenant</b>	<p>To choose a tenant:</p> <ol style="list-style-type: none"> <li>Click <b>Select Tenant</b>. The <b>Select Tenant</b> dialog box appears.</li> <li>From the <b>Select Tenant</b> dialog, click to choose a tenant in the left column. <ul style="list-style-type: none"> <li><b>Note</b> Beginning in Release 5.0(2), you can create contracts in the infra tenant. You can also export contracts from and import contracts to the infra tenant for shared services use cases.</li> </ul> </li> <li>Click <b>Select</b>. You return to the <b>Create Contract</b> dialog box.</li> </ol>
<b>Description</b>	Enter a description of the contract.
<b>Settings</b>	
<b>Scope</b>	<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>For inter-tenant communication, you will first create a contract with the <b>Global</b> scope in one of the tenants (for example, <b>tenant1</b>). This tenant's EPG will always be the provider of this contract.</p> <p>This contract will then be exported to the other tenant (for example, <b>tenant2</b>). For the other tenant that imports this contract, its EPG will be the consumer of the imported contract. If you want <b>tenant2</b>'s EPG to be the provider and <b>tenant1</b>'s EPG to be the consumer, then create a contract in <b>tenant2</b> and then export it to <b>tenant1</b>.</p>
<b>Add Filter</b>	<p>To choose a filter:</p> <ol style="list-style-type: none"> <li>Click <b>Add Filter</b>. The filter row appears with a <b>Select Filter</b> option.</li> <li>Click <b>Select Filter</b>. The <b>Select Filter</b> dialog box appears.</li> <li>From the <b>Select Filter</b> dialog, click to choose a filter in the left column then click <b>Select</b>. You return to the <b>Create Contract</b> dialog box.</li> </ol>

**Step 5** Click **Save** when finished.

**Step 6** Export the contract that you just created to another tenant.

For example, assume the following:

- The contract that you created in the procedure above is named **contract1** in tenant **tenant1**.
- The contract that you want to export is named **exported\_contract1** and you are exporting it to tenant **tenant2**.

- a) Navigate to the Contracts page (**Application Management > Contracts**).  
The configured contracts are listed.
- b) Select the contract that you just created.  
For example, scroll through the list until you see the contract **contract1** and click the box next to it to select it.
- c) Go to **Actions > Export Contract**.  
The **Export Contract** window appears.
- d) Click **Select Tenant**.  
The **Select Tenant** window appears.
- e) Select the tenant that you want to export the contract to, then click **Save**.  
For example, **tenant2**. You are returned to the **Export Contract** window.
- f) In the **Name** field, enter a name for the exported contract.  
For example, **exported\_contract1**.
- g) In the **Description** field, enter a description for the exported contract, if necessary.
- h) Click **Save**.  
The list of contracts appears again.

**Step 7** Configure the first tenant's EPG as the provider EPG, with the original contract, as the first part of the EPG communication configuration.

- a) Click the **Intent** button, then choose **EPG Communication**.  
The **EPG Communication** window appears.
- b) Click **Let's Get Started**.
- c) In the **Contract** area, click **Select Contract**.  
The **Select Contract** window appears.
- d) Locate and select the contract that you created at the beginning of these procedures.  
In this example, you would locate and select **contract1**.
- e) Click **Select**.  
The **EPG Communication** window appears.
- f) In the **Provider EPGs** area, click **Add Provider EPGs**.  
The **Select Provider EPGs** window appears.
- g) Leave the **Keep selected items** box checked, then select the first tenant's (**tenant1**) EPG.
- h) Click **Select**.  
The **EPG Communication** window appears.
- i) Click **Save**.

**Step 8** Configure the second tenant's EPG as the consumer EPG, with the exported contract, as the second part of the EPG communication configuration.

- a) Click the **Intent** button, then choose **EPG Communication**.

The **EPG Communication** window appears.

- b) Click **Let's Get Started**.
- c) In the **Contract** area, click **Select Contract**.

The **Select Contract** window appears.

- d) Locate and select the contract that you created at the beginning of these procedures.

In this example, you would locate and select **exported\_contract1**.

- e) Click **Select**.

The **EPG Communication** window appears.

- f) In the **Consumer EPGs** area, click **Add Consumer EPGs**.

The **Select Consumer EPGs** window appears.

- g) Leave the **Keep selected items** box checked, then select the second tenant's (**tenant2**) EPG.
- h) Click **Select**.

The **EPG Communication** window appears.

- i) Click **Save**.

---

## 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:

- a) Click **Select Contract**. The **Select Contract** dialog appears.
- b) 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:

- a) Click **Add Consumer EPGs**. The **Select Consumer EPGs** dialog appears.

**Note** EPGs within the tenant (where the contract is created) are displayed.

- b) 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:

- a) Click **Add Provider EPGs**. The **Select Provider EPGs** dialog appears.

**Note** EPGs within the tenant (where the contract is created) are displayed.

- b) 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.

**Note** If the chosen contract is an Imported Contract, the provider EPG selection is disabled.

- c) When finished, click **Select**. The **Select Provider EPGs** dialog box closes, and you return to the **EPG Communication Configuration** window.  
d) 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** 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 Cloud Context Profile**. The **Create Cloud Context Profile** dialog box appears.

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

**Table 12: Create Cloud Context Profile Dialog Box Fields**

Properties	Description
<b>Name</b>	Enter the name of the cloud context profile.
<b>Tenant</b>	To choose a tenant: <b>a.</b> Click <b>Select Tenant</b> . The <b>Select Tenant</b> dialog box appears. <b>b.</b> From the <b>Select Tenant</b> dialog, click to choose a tenant in the left column then click <b>Select</b> . You return to the <b>Create Cloud Context Profile</b> dialog box.
<b>Description</b>	Enter a description of the cloud context profile.
<b>Settings</b>	

Properties	Description
<b>Region</b>	<p>To choose a region:</p> <ol style="list-style-type: none"> <li>a. Click <b>Select Region</b>. The <b>Select Region</b> dialog box appears.</li> <li>b. From the <b>Select Region</b> dialog, click to choose a region in the left column then click <b>Select</b>. You return to the <b>Create Cloud Context Profile</b> dialog box.</li> </ol>
<b>VRF</b>	<p>To choose a VRF:</p> <ol style="list-style-type: none"> <li>a. Click <b>Select VRF</b>. The <b>Select VRF</b> dialog box appears.</li> <li>b. From the <b>Select VRF</b> dialog box, click to choose a VRF in the left column then click <b>Select</b>. You return to the <b>Create Cloud Context Profile</b> dialog box.</li> </ol>
<b>Add CIDR</b>	<p><b>Note</b> The following subnet is reserved and should not be used in this <b>Add CIDR</b> field: 192.168.100.0/24 (reserved by the CCR for the bridge domain interface)</p> <p><b>Note</b> You cannot add, delete, or edit a CIDR when VNet peering is enabled. You must disable VNet peering before adding, deleting or editing a CIDR. To disable VNet peering:</p> <ul style="list-style-type: none"> <li>• For the infra tenant, disable the <b>Hub Network Peering</b> option in the cloud context profile</li> <li>• For a user (non-infra) tenant, disable the <b>VNet Peering</b> option in the cloud context profile</li> </ul> <p>Enable VNet peering again after you have made the changes to the CIDR configuration.</p> <p><b>Note</b> Beginning in Release 5.0(2), you can add additional secondary CIDRs and subnets for infra VPCs (cloudCtxProfiles created by the cloud template). You cannot add primary CIDRs or modify the existing CIDRs created by the cloud template. After subnets are created under the user-created CIDRs, the subnets will be implicitly mapped to the overlay-2 VRF.</p> <p>To add a CIDR:</p> <ol style="list-style-type: none"> <li>a. Click <b>Add CIDR</b>. The <b>Add CIDR</b> dialog box appears.</li> <li>b. Enter the address in the <b>Address</b> field.</li> <li>c. Click <b>Add Subnet</b> and enter the subnet address in the <b>Address</b> field.</li> <li>d. Click to check (enabled) or uncheck (disabled) the <b>Primary</b> check box.</li> <li>e. When finished, click <b>Add</b>.</li> </ol>
<b>VNet Gateway Router</b>	<p>Click to check (enable) or uncheck (disable) in the <b>VNet Gateway Router</b> check box.</p>

Properties	Description
VNet Peering	Click to check (enable) or uncheck (disable) the Azure VNet peering feature. For more information on the VNet peering feature, see the <i>Configuring VNet Peering for Cloud APIC for Azure</i> document in the <a href="#">Cisco Cloud APIC documentation page</a> .

**Step 5** Click **Save** when finished.

## Configuring Virtual Machines in Azure

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

This topic provides the requirements for configuring the virtual machines in Azure. You can use these requirements to configure the virtual machines in Azure either before you configure the endpoint selectors for Cisco Cloud APIC or afterward. For example, you might go to your account in Azure and create a custom tag or label in Azure 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 Azure and create a custom tag or label in Azure afterward.

### Before you begin

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

**Step 1** Review your cloud context profile configuration to get the following information:

- VRF name
- Subnet information
- Subscription Id
- The resource group that corresponds to where the cloud context profile is deployed.

**Note** In addition to the information above, if you are using tag-based EPGs, you also need to know the tag names. The tag names are not available in the cloud context profile configuration.

To obtain the cloud context profile configuration information:

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 Azure virtual machine 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 Azure virtual machine.

**Step 2** Log in to the Azure portal account for the Cisco Cloud APIC user tenant and begin creating an Azure VM using the information you gathered from the cloud context profile configuration.

**Note** For information about how to create the VM in the Azure portal, see the Microsoft Azure documentation.

## 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 13: Create Backup Configuration Dialog Box Fields**

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

Properties	Description
Backup Object	



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

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

**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 14: Create Tech Support Dialog Box Fields**

Properties	Description
<b>General</b>	
<b>Name</b>	Enter the name of the tech support policy.
<b>Description</b>	Enter a description of the tech support.
<b>Settings</b>	

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

This section explains how to create a scheduler, which would be in User Laptop Browser local time and will be converted to the Cisco Cloud APIC default UTC time.

**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 Scheduler** dialog box appears.

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

**Table 15: Create Scheduler Dialog Box Fields**

Properties	Description
<b>General</b>	
<b>Name</b>	Enter the name of the trigger scheduler policy.
<b>Description</b>	Enter a description of the trigger scheduler.
<b>Settings</b>	

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

**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 16: Create Remote Location Dialog Box Fields**

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

**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 17: Create Login Domain Dialog Box Fields**

Properties	Description
<b>Name</b>	Enter the name of the login domain.
<b>Description</b>	Enter a description of the login domain.
<b>Realm</b>	Choose a realm: <ul style="list-style-type: none"> <li>• <b>Local</b></li> <li>• <b>LDAP</b>—Requires adding providers and choosing an authentication type.</li> <li>• <b>RADIUS</b>—Requires adding providers.</li> <li>• <b>TACACS+</b>—Requires adding providers.</li> <li>• <b>SAML</b>—Requires adding providers.</li> </ul>
<b>Providers</b>	To add a provider: <ol style="list-style-type: none"> <li>a. Click <b>Add Providers</b>. The <b>Select Providers</b> dialog appears with a list of providers in the left pane.</li> <li>b. Click to choose a provider.</li> <li>c. Click <b>Select</b> to add the provider.</li> </ol>
<b>Advanced Settings</b>	Displays the <b>Authentication Type</b> and <b>LDAP Group Map Rules</b> fields.

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

**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 18: Create Role Dialog Box Fields**

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



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"> <li>• <b>aaa</b>—Used for configuring authentication, authorization, accounting and import/export policies.</li> <li>• <b>access-connectivity-11</b>—Used for Layer 1 configuration under infra. Example: selectors and port Layer 1 policy configurations.</li> <li>• <b>access-connectivity-12</b>—Used for Layer 2 configuration under infra. Example: Encap configurations on selectors, and attachable entity.</li> <li>• <b>access-connectivity-13</b>—Used for Layer 3 configuration under infra and static route configurations under a tenant's L3Out.</li> <li>• <b>access-connectivity-mgmt</b>—Used for management infra policies.</li> <li>• <b>access-connectivity-util</b>—Used for tenant ERSPAN policies.</li> <li>• <b>access-equipment</b>—Used for access port configuration.</li> <li>• <b>access-protocol-11</b>—Used for Layer 1 protocol configurations under infra.</li> <li>• <b>access-protocol-12</b>—Used for Layer 2 protocol configurations under infra.</li> <li>• <b>access-protocol-13</b>—Used for Layer 3 protocol configurations under infra.</li> <li>• <b>access-protocol-mgmt</b>—Used for fabric-wide policies for NTP, SNMP, DNS, and image management.</li> <li>• <b>access-protocol-ops</b>—Used for operations-related access policies such as cluster policy and firmware policies.</li> <li>• <b>access-protocol-util</b>—Used for tenant ERSPAN policies.</li> <li>• <b>access-qos</b>—Used for changing CoPP and QoS-related policies.</li> <li>• <b>admin</b>—Complete access to everything (combine ALL roles)</li> <li>• <b>fabric-connectivity-11</b>—Used for Layer 1 configuration under the fabric. Example: selectors and port Layer 1 policy and VNET protection.</li> </ul>

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

Properties	Description
	<ul style="list-style-type: none"> <li>• <b>ops</b>—Used for operational policies including monitoring and troubleshooting policies such as atomic counter, SPAN, TSW, tech support, traceroute, analytics, and core policies.</li> <li>• <b>tenant-connectivity-11</b>—Used for Layer 1 connectivity changes, including bridge domains and subnets.</li> <li>• <b>tenant-connectivity-12</b>—Used for Layer 2 connectivity changes, including bridge domains and subnets.</li> <li>• <b>tenant-connectivity-13</b>—Used for Layer 3 connectivity changes, including VRFs.</li> <li>• <b>tenant-connectivity-mgmt</b>—Used for tenant in-band and out-of-band management connectivity configurations and for debugging/monitoring policies such as atomic counters and health score.</li> <li>• <b>tenant-connectivity-util</b>—Used for atomic counter, diagnostic, and image management policies on leaf switches and spine switches.</li> <li>• <b>tenant-epg</b>—Used for managing tenant configurations such as deleting/creating endpoint groups, VRFs, and bridge domains.</li> <li>• <b>tenant-ext-connectivity-12</b>—Used for managing tenant L2Out configurations.</li> <li>• <b>tenant-ext-connectivity-13</b>—Used for managing tenant L3Out configurations.</li> <li>• <b>tenant-ext-connectivity-mgmt</b>—Used as write access for firmware policies.</li> <li>• <b>tenant-ext-connectivity-util</b>—Used for debugging/monitoring/observer policies such as traceroute, ping, oam, and eptrk.</li> <li>• <b>tenant-ext-protocol-11</b>—Used for managing tenant external Layer 1 protocols. Generally only used for write access for firmware policies.</li> <li>• <b>tenant-ext-protocol-12</b>—Used for managing tenant external Layer 2 protocols. Generally only used for write access for firmware policies.</li> <li>• <b>tenant-ext-protocol-13</b>—Used for managing tenant external Layer 3 protocols such as BGP, OSPF, PIM, and IGMP.</li> <li>• <b>tenant-ext-protocol-mgmt</b>—Used as write access for firmware policies.</li> </ul>

Properties	Description
	<ul style="list-style-type: none"> <li>• <b>tenant-ext-protocol-util</b>—Used for debugging/monitoring/observer policies such as traceroute, ping, oam, and eptrk.</li> <li>• <b>tenant-network-profile</b>—Used for managing tenant configurations, such as deleting and creating network profiles, and deleting and creating endpoint groups.</li> <li>• <b>tenant-protocol-l1</b>—Used for managing configurations for Layer 1 protocols under a tenant.</li> <li>• <b>tenant-protocol-l2</b>—Used for managing configurations for Layer 2 protocols under a tenant.</li> <li>• <b>tenant-protocol-l3</b>—Used for managing configurations for Layer 3 protocols under a tenant.</li> <li>• <b>tenant-protocol-mgmt</b>—Only used as write access for firmware policies.</li> <li>• <b>tenant-protocol-ops</b>—Used for tenant traceroute policies.</li> <li>• <b>tenant-protocol-util</b>—Used for debugging/monitoring/observer policies such as traceroute, ping, oam, and eptrk.</li> <li>• <b>tenant-qos</b>—Only used as Write access for firmware policies.</li> <li>• <b>tenant-security</b>—Used for Contract related configurations for a tenant.</li> <li>• <b>vmm-connectivity</b>—Used to read all the objects in APIC's VMM inventory required for VM connectivity.</li> <li>• <b>vmm-ep</b>—Used to read VM and Hypervisor endpoints in the APIC's VMM inventory.</li> <li>• <b>vmm-policy</b>—Used for managing policies for VM networking.</li> <li>• <b>vmm-protocol-ops</b>—Not used by VMM policies.</li> <li>• <b>vmm-security</b>—Used for Contract related configurations for a tenant.</li> </ul>

**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:  
a) Click **Select Security Domain**. The **Select Security Domain** dialog box appears.  
b) 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 19: Create Certificate Authority Dialog Box Fields**

Properties	Description
<b>Name</b>	Enter the name of the certificate authority.
<b>Description</b>	Enter a description of the certificate authority.

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

**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 20: Create Key Ring Dialog Box Fields**

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



Properties	Description
<b>Modulus</b>	Click the <b>Modulus</b> drop-down list to choose from the following: <ul style="list-style-type: none"> <li>• <b>MOD 512</b></li> <li>• <b>MOD 1024</b></li> <li>• <b>MOD 1536</b></li> <li>• <b>MOD 2048</b>—(Default)</li> </ul>
<b>Certificate</b>	Enter the certificate information in the <b>Certificate</b> 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 21: Create Local User Dialog Box Fields**

Properties	Description
<b>Name</b>	Enter the username of the local user.
<b>Password</b>	Enter the password for the local user.
<b>Confirm Password</b>	Reenter the password for the local user.
<b>Description</b>	Enter a description of the local user.
<b>Settings</b>	
<b>Account Status</b>	To choose the account status: <ul style="list-style-type: none"> <li>• <b>Active</b>—Activates the local user account.</li> <li>• <b>Inactive</b>—Deactivates the local user account.</li> </ul>
<b>First Name</b>	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"> <li>a. Click <b>Add Security Domain</b>. The <b>Add Security Domain</b> dialog box appears.</li> <li>b. Click <b>Select Security Domain</b>. The <b>Select Security Domain</b> dialog box appears with a list of security domains in the left pane.</li> <li>c. Click to choose a security domain.</li> <li>d. Click <b>Select</b> to add the security domain. You return to the <b>Add Security Domain</b> dialog box.</li> <li>e. Add a user role: <ol style="list-style-type: none"> <li>1. From the <b>Add Security Domain</b> dialog box, click <b>Select Role</b>. The <b>Select Role</b> dialog box appears with a list of roles in the left pane.</li> <li>2. Click to choose a role.</li> <li>3. Click <b>Select</b> to add the the role. You return to the <b>Add Security Domain</b> dialog box.</li> <li>4. From the <b>Add Security Domain</b> dialog box, click the <b>Privilege Type</b> drop-down list and choose <b>Read Privilege</b> or <b>Write Privilege</b>.</li> <li>5. Click the check mark on the right side of the <b>Privilege Type</b> drop-down list to confirm.</li> <li>6. Click <b>Add</b> when finished. You return to the <b>Create Local User</b> dialog box where you can add another security domain.</li> </ol> </li> </ol>

**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 22: Create Local User Dialog Box Fields: Advanced Settings**

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

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

**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 24](#).

**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 **cAPIC Setup**. The **Set up - Overview** dialog box appears with options for **DNS and NTP Servers**, **Region Management**, and **Smart Licensing**.

- Step 4** For **Region Management**, click **Edit Configuration**. The **Set Up - Region Management** dialog box appears with a list of managed regions.
- Step 5** 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. The **Cloud Routers** and **Inter-Site Connectivity** check boxes are enabled.
- Step 6** To deploy cloud routers locally to this region, click to place a check mark in the **Cloud Routers** check box.
- Step 7** To enable the cloud routers in the region to connect to on-premises ACI sites, click to place a check mark in the **Inter-Site Connectivity** check box. The **Cloud Routers** check box is automatically checked.
- Step 8** To configure the fabric infra connectivity for the cloud site, click **Next**.
- Step 9** Add the Fabric Autonomous System number for the Azure Cloud Site.
- Step 10** To specify the subnet, click **Add Subnet 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 11** 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 12** Enter a username in the **Username** text box.
- Step 13** Enter a password in the **Password** and **Confirm Password** text boxes.
- Step 14** 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.
- Step 15** (Optional) To specify the license token, enter the product instance registration token in the **License Token** text box.
- Note** If no token is entered, the CSR will be in EVAL mode.
- Step 16** To configure inter-site connectivity, click **Next**.
- Step 17** 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 18** Enter the OSPF area ID in the **OSPF Area Id** text box.
- Step 19** To add an external subnet pool, click **Add External Subnet** and enter a subnet pool in the text box.
- Step 20** When you have configured all the connectivity options, click **Next** at the bottom of the page.
- The **Cloud Resource Naming Rules** page appears, which is described in detail in the [Cloud Resources Naming, on page 73](#) section. If you don't need to make any changes to the naming rules, you can skip this page.
- Step 21** Click **Save and Continue** when finished.

## Configuring Smart Licensing

This task demonstrates how to set up smart licensing in the Cisco Cloud APIC.

### Before you begin

You need the product instance registration token.

- 
- 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 **Set Up cAPIC**. The **Set up - Overview** dialog box appears with options for **DNS Servers**, **Region Management**, and **Smart Licensing**.
- Step 4** To register the Cloud APIC to Cisco's unified license management system: From **Smart Licensing**, click **Register**. The **Smart Licensing** dialog appears.
- Step 5** Choose a transport setting:
- **Direct to connect to Cisco Smart Software Manager (CSSM)**
  - **Transport Gateway/Smart Software Manager Satellite**
  - **HTTP/HTTPS Proxy**
- Note** An IP address is also required when choosing **HTTP/HTTPS Proxy**.
- Step 6** Enter the product instance registration token in the provided text box.
- Step 7** Click **Register** when finished.
- 

## Cloud Resources Naming

Prior to Cloud APIC Release 5.0(2), the cloud resources created by the Cloud APIC in Azure were assigned names that were derived from the names of the ACI objects:

- Resource groups were created based on the Tenant, VRF, and region. For example, `CAPIC_<tenant>_<vrf>_<region>`.
- VNET names matched the name of the Cloud APIC VRF.
- Subnet names were derived from the CIDR address space. For example, `subnet-10.10.10.0_24` for the `10.10.10.0/24` cloud subnet.
- The cloud application name was derived from the EPG name and the application profile name. For example, `<epg-name>_cloudapp_<app-profile-name>`

This approach is not ideal for deployments with strict cloud resource naming conventions and it does not follow the Azure best practices for naming and tagging of cloud resources.

Starting with Cloud APIC Release 5.0(2), you can create a global naming policy on the Cloud APIC, which allows you to define a custom cloud resources naming convention for all objects deployed from the Cloud APIC into the Azure cloud. You can define custom naming rules for all cloud resources during the first time setup wizard of the Cloud APIC, with the exception of the **Resource group** name used for the Cloud APIC ARM template deployment. The resource group name for the template is defined when you first deploy it and cannot be changed after. In addition to the global policy, you can also explicitly define the names of the cloud resources created from each Cloud APIC object using the REST API.



**Note** Keep in mind that even with custom naming policy, once a cloud resource is created, you will not be able to modify the name. If you want to change the name of an existing cloud resource, you would need to delete all configured cloud resources and recreate them. Cloud resources to be deleted include overlay-2 CIDR and subnets, Cisco Cloud Services Router 1000Vs deployed by Cloud APIC and therefore IPsec tunnels from the CSRs to every remote site.

## Variables Available for Naming Rules

When creating your cloud resources naming policy, you can use the following variables to dynamically define the name of the cloud resource based on the Cloud APIC objects:

- `${tenant}` – the resource will include the name of the Tenant
- `${ctx}` – the resource will include the name of the VRF
- `${ctxprofile}` – the resources will include the cloud context profile, which is a VRF deployed in a given cloud region
- `${app}` – the resource will include the name of the application profile.
- `${epg}` – the resource will include the name of the EPG.
- `${contract}` – the resource will include the name of the contract
- `${region}` – the resource will include the name of the cloud region
- `${priority}` – the resource will include the name of the network security group (NSG) rule priority. This number is allocated automatically to ensure that each NSG rule name is unique

When you define a global naming policy using one or more of the above variables, Cloud APIC validates the string to ensure that all mandatory variables are present and no invalid string is specified.

There is a maximum name length limit in Azure. If the length of the name exceeds the length supported by the cloud provider, it rejects the config and Cloud APIC raises a fault that the resource creation failed. You can then check the fault for details and correct the naming rules. The maximum length limits at the time of Cloud APIC, Release 5.0(2) are listed below, for the latest up-to-date information and any changes to the length limit, consult the Azure documentation.

The following table provides a summary of which cloud resources support each of the naming variables above. Cells denoted with an asterisk (\*) indicate variables that are mandatory for that type of cloud resource. Cells denoted with a plus sign (+) indicate that at least one of these variables is mandatory for that type of cloud resource; for example, for VNET resources you can provide `${ctx}`, or `${ctxprofile}`, or both.

**Table 23: Supported Variables for Cloud Resources**

Azure Resource	<code>\${tenant}</code>	<code>\${ctx}</code>	<code>\${ctxprofile}</code>	<code>\${subnet}</code>	<code>\${app}</code>	<code>\${epg}</code>	<code>\${contract}</code>	<code>\${region}</code>	<code>\${priority}</code>
Resource Group	Yes*	Yes*						Yes*	
Max Length: 90									

Azure Resource	\$(tenant)	\$(ctx)	\$(cbprofile)	\$(subnet)	\$(app)	\$(epg)	\$(contract)	\$(region)	\$(priority)
Virtual Network (VNET) Max Length: 64	Yes	Yes+	Yes+					Yes	
Subnet Max Length: 80	Yes	Yes	Yes	Yes*				Yes	
Application Security Group (ASG) Max Length: 80	Yes				Yes*	Yes*		Yes	
Network Security Group (NSG) Max Length: 80	Yes				Yes*	Yes*		Yes	
Network Security Group Rule Max Length: 80	Yes						Yes		Yes* (auto)

## Naming Rules Guidelines and Limitations

When configuring custom rules for naming cloud resources, the following restrictions apply:

- You define global naming policy during the Cloud APIC's first time setup using two sets of naming rules:
  - Hub Resource Naming Rules** define names for the Hub Resource Group, Hub VNET, and Overlay-1 CIDR subnet in the Infra Tenant, as well as the subnet prefixes for subnets that are created automatically by the system in the Infra tenant.
  - Cloud Resource Naming Rules** define the names of the Network Security Group (NSG), Application Security Group (ASG), and subnets you create in the Infra Tenant, as well as the names of all resources (Resource Groups, Virtual Networks, Subnets, NSG, ASG) in user Tenants.

After you define the naming rules, you will be required to review and confirm them. Keep in mind that you must confirm the naming rules before any cloud resources are deployed.

- Once a cloud resource is created, its name cannot be changed and the naming policy cannot be updated in the GUI. If you upgrade your Cloud APIC to Release 5.0(2) with some resources already deployed in Azure, you will also not be able to change the global custom naming rules.

If you want to change the names of the existing cloud resources or the policy, you would need to delete the deployed resources before being able to update the global naming policy in the GUI.

In these cases you can use the REST API to explicitly assign custom names to any new resources you create.

- When updating cloud resources naming via REST API, we recommend you do not import configuration at the same time.

We recommend you define any naming rules first. Then any tenant configuration.

We recommend that you do not change the naming policy after the tenant configuration is deployed.

## Viewing Cloud Resource Naming Rules

You initially define the cloud resource naming rules in the Region Management part of the first time setup wizard when you deploy your Cloud APIC, which is described in the *Cisco Cloud APIC Installation Guide*. After the initial setup, you can view the rules you configured in the **System Configuration** screen of your Cloud APIC GUI as described in this section.

Note that the information in this screen is presented in read-only view and if you want to change the rules any time after the original deployment, you will need to re-run the first time setup wizard .

**Step 1** Log in to your Cloud APIC GUI.

**Step 2** Navigate to the **Cloud Resource Naming Rules** screen.

The screenshot displays the Cisco Cloud APIC System Configuration interface. The left sidebar shows the navigation menu with 'Infrastructure' expanded and 'System Configuration' selected. The main content area is titled 'System Configuration' and has tabs for 'General', 'Management Access', 'Cloud Resource Naming Rules', 'Controllers', and 'Event Analytics'. A warning banner at the top states: 'Please go to cAPIC Setup Region Management to manage Hub and Cloud Resource Naming Rules. Go to cAPIC Setup'. Below this, a diagram illustrates the process: 'Create the Cloud APIC policy' (policyName) leads to 'Mapped Cloud Resource' (Cloud Resource 1 and 2) and 'Naming Rule' (\$[Policy]\_resource-1 and -2), which then result in 'Cloud resources on Azure get created with the generated names based on the rules from Cloud APIC' (Cloud Resource 1 and 2 with policyName\_resource-1 and -2). Below the diagram are two tables:

Managed Region	Resource Group Name	Virtual Network Name	Subnet Name Prefix	Cloud Subnet Example
Canada Central	JMR1-1	overlay-1	subnet-	subnet-1.1.1.1_28
Central US	CAPIC_infra_overlay-1_centralus	overlay-1	subnet-	subnet-1.1.1.1_28

Cloud Resource	Mapped ACI Object	Naming Rule	Cloud Resource Example

- In the **Navigation** sidebar, expand the **Infrastructure** category.
- From the **Infrastructure** category, select **System Configuration**.
- In the **System Configuration** screen, select the **Cloud Resource Naming Rules** tab.

In the **Cloud Resource Naming Rules** tab, you can see a summary of the currently configured rules for the names of resources that you deploy in the cloud site from your Cloud APIC.

If you did not configure custom naming rules before, the default rules are listed here, which use the Cloud APIC object names for cloud resources.

If you have not accepted the naming rules you have defined during the first time setup, a warning banner will be displayed across the top of the screen.



**Note** Keep in mind that you must confirm the naming rules before any cloud resources are deployed.

# Configuring Cisco Cloud APIC Using the REST API

## Creating a Tenant Using the REST API

There are two types of subscriptions: own and shared. Each subscription type has a primary tenant. You choose the own subscription when creating a new managed or unmanaged tenant. You choose the shared subscription when creating a tenant that inherits the managed or unmanaged settings of an existing primary tenant. This section demonstrates how to create a managed and unmanaged tenant with the own type of subscription and how to create a shared subscription.

This section demonstrates how to create a tenant using the REST API using sample POST requests from the body of Postman.

### Step 1 Create an own subscription.

- a) To create an unmanaged tenant using a client secret:

```
POST https://<cloud-apic-ip-address>/api/mo/uni.xml

<fvTenant name="{{primary-tenant-name}}">
  <cloudAccount id="{{user-tenant-subscription-id}}" vendor="azure" accessType="credentials"
status="">
    <cloudRsCredentials tDn="uni/tn-{{primary-tenant-name }}/credentials-{{ primary-tenant-name
}}"/>
  </cloudAccount>
  <cloudCredentials name="{{ primary-tenant-name }}" keyId="{{application_key_id}}"
key="{{client_secret_key}}">
    <cloudRsAD tDn="uni/tn-{{ primary-tenant-name }}/ad-{{active_directory_id}}"/>
  </cloudCredentials>
  <cloudAD name="{{active_directory_name}}" id="{{active_directory_id}}" />
  <fvRsCloudAccount tDn="uni/tn-{{ primary-tenant-name }}/act-[[{ user-tenant-subscription-id
}]]-vendor-azure" status="" />
</fvTenant>
```

- b) To create a managed tenant:

```
POST https://<cloud-apic-ip-address>/api/mo/uni.xml

<fvTenant name="{{ primary-tenant-name }}">
  <cloudAccount id="{{ user-tenant-subscription-id }}" vendor="azure" accessType="managed"
status="" />
  <fvRsCloudAccount tDn="uni/tn-{{ primary-tenant-name }}/act-[[{ user-tenant-subscription-id
}]]-vendor-azure" status="" />
</fvTenant>
```

### Step 2 Create a shared subscription:

```
POST https://<cloud-apic-ip-address>/api/mo/uni.xml
```

```
<fvTenant name="{{ primary-tenant-name }}">
  <fvRsCloudAccount tDn="uni/tn-{{ primary-tenant-name }}/act-[[{ user-tenant-subscription-id
  }}]-vendor-azure" status=""/>
</fvTenant>
```

## 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.

To create a cloud context profile:

#### Example:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
  <fvTenant name="tn15">

    <cloudCtxProfile name="cProfilewestus15">
      <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-azure/region-westus"/>
      <cloudRsToCtx tnFvCtxName="ctx151"/>
      <cloudCidr addr="15.151.0.0/16" primary="true" status="">
```

```

        <cloudSubnet ip="15.151.1.0/24" name="GatewaySubnet" usage="gateway">
            <cloudRsZoneAttach tDn="uni/clouddomp/provp-azure/region-westus/zone-default"/>
        </cloudSubnet>
        <cloudSubnet ip="15.151.2.0/24" name="albsubnet" >
            <cloudRsZoneAttach tDn="uni/clouddomp/provp-azure/region-westus/zone-default"/>
        </cloudSubnet>
        <cloudSubnet ip="15.151.3.0/24" name="subnet" usage="">
            <cloudRsZoneAttach tDn="uni/clouddomp/provp-azure/region-westus/zone-default"/>
        </cloudSubnet>
    </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:

```

<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
    <cloudDomP name="default">
        <cloudProvP vendor="azure">
            <cloudRegion adminSt="managed" name="eastus"><cloudZone name="default"/></cloudRegion>
            <cloudRegion adminSt="managed" name="eastus2"><cloudZone name="default"/></cloudRegion>
            <cloudRegion adminSt="managed" name="westus"><cloudZone name="default"/></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
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
    <fvTenant name="t15">
        <vzFilter name="rule1">
            <vzEntry etherT="ip" dToPort="22" prot="tcp" dFromPort="22" name="ssh"/>
            <vzEntry etherT="ip" prot="unspecified" name="any"/>
        </vzFilter>
        <vzFilter name="rule2">
            <vzEntry etherT="ip" dToPort="http" prot="tcp" dFromPort="http" name="http"/>
        </vzFilter>
    </fvTenant>
</polUni>

```

```

    </vzFilter>
    <vzFilter name="rule3">
      <vzEntry etherT="ip" dToPort="22" prot="tcp" dFromPort="22" name="ssh"/>
    </vzFilter>
    <vzFilter name='all_rule'>
      <vzEntry etherT="ip" prot="unspecified" name="any"/>
    </vzFilter>

    <vzBrCP name="c1">
      <vzSubj name="c1">
        <vzRsSubjFiltAtt tnVzFilterName="rule2"/>
        <vzRsSubjGraphAtt tnVnsAbsGraphName="c13_g1"/>
        <vzRsSubjFiltAtt tnVzFilterName="rule3"/>
        <vzRsSubjFiltAtt tnVzFilterName="all_rule"/>
      </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
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
  <fvTenant name="tn15">
    <fvRsCloudAccount tDn="uni/tn-infra/act-[<subscription id>]-vendor-azure" />

    <fvCtx name="ctx151"/>

    <cloudVpnGwPol name="VgwPol1"/>
    <cloudApp name="a1">

  </cloudApp>

</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:**

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
  <fvTenant name="tn15">
    <fvRsCloudAccount tDn="uni/tn-infra/act-[<subscription id>]-vendor-azure" />

    <fvCtx name="ctx151"/>

    <cloudVpnGwPol name="VgwPol1"/>
    <cloudApp name="a1">

      <cloudEPg name="epg1">
        <cloudRsCloudEPgCtx tnFvCtxName="ctx151"/>
        <cloudEPSelector matchExpression="custom:tag1=='value1'" name="selector-1"/>
      </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:**

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
  <fvTenant name="tn15">
    <fvRsCloudAccount tDn="uni/tn-infra/act-[<subscription id>]-vendor-azure" />

    <fvCtx name="ctx151"/>

    <cloudVpnGwPol name="VgwPol1"/>
    <cloudApp name="a1">
      <cloudExtEPg routeReachability="internet" name="extEpg-1">
    </cloudExtEPg>
    </cloudApp>
  </fvTenant>
</polUni>
```

```

        <fvRsCons tnVzBrCPName="extEpg-1"/>
        <cloudRsCloudEPgCtx tnFvCtxName="ctx151"/>
        <cloudExtEPSelector name="extSelector1" subnet="0.0.0.0/0"/>
    </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 24](#).

### Before you begin

To create a cloud template:

```

<polUni>
  <fvTenant name="infra">
    <cloudtemplateInfraNetwork name="default" numRemoteSiteSubnetPool="2" numRoutersPerRegion="2"
status="" vrfName="overlay-1">
      <cloudtemplateProfile name="default" routerPassword="cisco123" routerUsername="cisco"
routerThroughput="250M" routerLicenseToken="thisismysrtoken" />
    </cloudtemplateProfile>
    <cloudtemplateExtSubnetPool subnetpool="10.20.0.0/16"/>

    <cloudtemplateIntNetwork name="default">
      <cloudRegionName provider="azure" region="westus"/>
      <cloudRegionName provider="azure" region="westus2"/>
    </cloudtemplateIntNetwork>

    <cloudtemplateExtNetwork name="default">
      <cloudRegionName provider="azure" region="westus2"/>

    <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>

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

```

## Defining Global Cloud Resource Naming Rules or Overriding Specific Object's Name

This section provides an example REST API POST you can use to configure a global policy for naming your cloud resources or override a specific cloud resource's name.



**Note** To ensure that any custom naming conventions can be supported, cloud resource names can be defined on a per-object basis. These explicit name overrides are not available in the Cloud APIC GUI and can be done using REST API only. We recommend using the global cloud resource naming policy to define the names. Explicit name overrides should be used only when naming requirements cannot be met using the global naming policy.

### Step 1 To create Hub Resource Naming Rules:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
  <fvTenant name="infra">
    <cloudtemplateInfraNetwork name="default" numRemoteSiteSubnetPool="2"
      numRoutersPerRegion="2" status="" vrfName="overlay-1">
      <cloudtemplateIntNetwork name="default">
        <cloudRegionName provider="azure" region="west's" status=""
          <cloudtemplateRegionNameCustomization ctxProfileName="infra-vnet"
            resourceGroupName="infra-rh" subnetNamePrefix="snet-" />
        </cloudRegionName>
      </cloudtemplateIntNetwork>
    </cloudtemplateInfraNetwork>
  </fvTenant>
</polUni>
```

### Step 2 To create Cloud Resource Naming Rules:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
  <cloudDomP name="default">
    <cloudNaming
      azResourceGroup="${tenant}-network-${ctx}-${region}-rg"
      azVirtualNetwork="${tenant}-${ctxprofile}-vnet"
      azSubnet="${tenant}-${ctxprofile}-snet-${subnet}"
      azNetworkSecurityGroup="${app}-${epg}-nsg"
      azApplicationSecurityGroup="${app}-${epg}-asg"
      azNetworkSecurityGroupRule="${contract}--${priority}"
      reviewed="yes" />
    </cloudDomP>
  </polUni>
```

### Step 3 To override an Azure cloud resource name corresponding to a specific Cloud APIC object:

You can use the same variables (for example, `${tenant}`) when specifying the custom name using the API.

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<fvTenant name="ExampleCorp" status="">
  <fvRsCloudAccount status="" tDn="uni/tn-infra/act-[<infra-subscription>]-vendor-azure"/>
  <fvCtx name="VRF1"/>
</fvTenant>
```

```
<cloudApp name="App1">
  <cloudEPg name="Db" azNetworkSecurityGroup="db-nsg" azApplicationSecurityGroup="db-asg-${region}">

    <cloudRsCloudEPgCtx tnFvCtxName="VRF1"/>
    <cloudEPSelector matchExpression="custom:EPG=='db'" name="100"/>
  </cloudEPg>
</cloudApp>
<cloudCtxProfile name="c02" azResourceGroup="custom-tc-rg1" azVirtualNetwork="vnet1">
  <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-azure/region-westus"/>
  <cloudRsToCtx tnFvCtxName="VRF1"/>
  <cloudCidr addr="10.20.20.0/24" name="cidr1" primary="yes" status="">
    <cloudSubnet ip="10.20.20.0/24" name="subnet1" azSubnet="s1" status="">
      <cloudRsZoneAttach status="" tDn="uni/clouddomp/provp-azure/region-westus/zone-default"/>
    </cloudSubnet>
  </cloudCidr>
</cloudCtxProfile>
</fvTenant>
```

---





## CHAPTER 5

# Viewing System Details

- [Viewing Application Management Details, on page 85](#)
- [Viewing Cloud Resource Details, on page 86](#)
- [Viewing Operations Details, on page 88](#)
- [Viewing Infrastructure Details, on page 90](#)
- [Viewing Administrative Details, on page 90](#)
- [Viewing Health Details Using the Cisco Cloud APIC GUI, on page 92](#)

## 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 24: 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.

Subtab Name	Description
Services	Contains the following two subtabs and information: <ul style="list-style-type: none"> <li>• <b>Devices</b>—Displays the devices as rows in a summary table.</li> <li>• <b>Service Graphs</b>—Displays service graphs as rows in a summary table.</li> </ul>
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 Name == 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.
- **Topology**—Provides visual relationship between an object and other related objects. The chosen object is displayed at the center.
- **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.

**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, VNET, router, security group (application security group/network security group), endpoint, VM, and cloud service.

Beginning with Release 5.0(2), for the **Endpoints** subtab, search based on *Cloud Tag* attribute is supported.

**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 25: Cloud Resource Subtabs**

Subtab Name	Description
<b>Regions</b>	Displays regions as rows in a summary table.
<b>Virtual Networks</b>	Displays VNETs as rows in a summary table.
<b>Routers</b>	Displays routers as rows in a summary table.
<b>Security Groups</b>	Displays security groups as rows in a summary table.
<b>Endpoints</b>	Displays endpoints as rows in a summary table.
<b>Virtual Machines</b>	Displays the VMs as rows in a summary table.
<b>Cloud Services</b>	Contains the following subtabs: <ul style="list-style-type: none"> <li>• <b>Cloud Service</b> Tab—Displays cloud services as rows in a summary table.</li> <li>• <b>Target Groups</b> Tab—Displays target groups as rows in a summary table.</li> </ul>

**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.

For the **Endpoints** subtab, you can narrow down the search based on a cloud tag, by entering a **key** or **value** term. If you want to search based on both terms, click the (+) displayed as a superscript to the **key** or **value** term (depending on which was entered first). Cloud tag filters cannot be edited. To modify a search, first delete the filters, and then enter the desired **key** or **value** term again. Search based on multiple cloud tag filters is supported.

**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. Beginning with Release 5.0(2), the cloud tags associated with endpoints are displayed.
- **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 26: Operations Subtabs**

Subtab Name	Description
<b>Event Analytics</b>	Contains the following subtabs: <ul style="list-style-type: none"> <li>• <b>Faults</b> Tab—Displays faults as rows in a summary table.</li> <li>• <b>Fault Records</b> Tab—Displays fault records as rows in a summary table.</li> <li>• <b>Events</b> Tab—Displays events as rows in a summary table.</li> <li>• <b>Audit Logs</b> Tab—Displays audit logs as rows in a summary table.</li> </ul>
<b>Active Sessions</b>	Displays a list of active users who are logged into Cloud APIC.

Subtab Name	Description
<b>Backup &amp; Restore</b>	Contains the following subtabs: <ul style="list-style-type: none"> <li>• <b>Backups</b> Tab—Displays backup as rows in a summary table.</li> <li>• <b>Backup Policies</b> Tab—Displays backup policies as rows in a summary table.</li> <li>• <b>Job Status</b> Tab—Displays the job status as rows in a summary table.</li> <li>• <b>Event Analytics</b> Tab—Contains the following subtabs:               <ul style="list-style-type: none"> <li>• <b>Faults</b> Tab—Displays faults as rows in a summary table.</li> <li>• <b>Events</b> Tab—Displays events as rows in a summary table.</li> <li>• <b>Audit Logs</b> Tab—Displays audit logs as rows in a summary table.</li> </ul> </li> </ul>
<b>Tech Support</b>	Contains the following subtabs: <ul style="list-style-type: none"> <li>• <b>Tech Support</b> Tab—Displays tech support policies as rows in a summary table.</li> <li>• <b>Core Logs</b> Tab—Displays core logs as rows in a summary table.</li> </ul>
<b>Firmware Management</b>	Contains the following subtabs: <ul style="list-style-type: none"> <li>• <b>General</b> Tab—Displays general firmware management information, such as Current Firmware Version, Upgrade Status.</li> <li>• <b>Images</b> Tab—Displays a list of images.</li> <li>• <b>Event Analytics</b> Tab—Contains the following subtabs:               <ul style="list-style-type: none"> <li>• <b>Faults</b> Tab—Displays faults as rows in a summary table.</li> <li>• <b>Events</b> Tab—Displays events as rows in a summary table.</li> <li>• <b>Audit Logs</b> Tab—Displays audit logs as rows in a summary table.</li> </ul> </li> </ul>
<b>Schedulers</b>	Displays scheduler policies as rows in a summary table.
<b>Remote Locations</b>	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 27: Infrastructure Subtabs**

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

**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 28: Administrative Subtabs**

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

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

**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.

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 user, choose User ID == admin (where admin is a user ID. ).

## 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 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 includes a 'System' section with several widgets:

- Health Summary:** A large orange banner indicating a 'Major' issue.
- 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' with a warning icon.
- Smart License Authorization Status:** Shows 'Evaluation' with a warning icon and '75 days remaining'.

Below the System section is a 'Cloud Resources Summary' section for Azure, showing counts for Regions (0 Total), Virtual Networks (2 Total), Routers (4 Total), Endpoints (0), and 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 Event Analytics window. The left sidebar is the same as in the previous screenshot. The main content area is titled 'Event Analytics' and includes a 'Faults' tab. The 'Severity' filter is set to 'Critical'. The table below shows the following data:

Acked	Severity	Code	Affected object	Description	Lifecycle	Creation Time	
<input type="checkbox"/>	No	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"/>	No	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 '10 Rows' dropdown and a 'Page 1 of 1' indicator.

**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/pool-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
No	Critical	F0104	topology/pool-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
No	Major	F3442	acct-[infra]/region-[eastus]/context-[overlay-1]-addr-[10.10.0.128/20]/csc-[ct_routerp_eastus_1_0]/nstoper	Operational State of the hcloud InstanceOper is down with [computeVirtualMachinesClientCreateOrUpdate: Failure sending request: StatusCode=404 -- Original Error: Code="ResourceGroupNotFound" Message="Resource group 'APIC-infra-mininet-fchazel-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]/csc-[ct_routerp_centralus_1_0]/nstoper	Operational State of the hcloud InstanceOper is down with [computeVirtualMachinesClientCreateOrUpdate: Failure sending request: StatusCode=404 -- Original Error: Code="ResourceGroupNotFound" Message="Resource group 'APIC-infra-mininet-fchazel-centralus' could not be found."]	raised	Sep 11 2019 07:38:27pm
No	Major	F3442	acct-[infra]/region-[eastus]/context-[overlay-1]-addr-[10.10.0.128/20]/csc-[ct_routerp_eastus_0_0]/nstoper	Operational State of the hcloud InstanceOper is down with [computeVirtualMachinesClientCreateOrUpdate: Failure sending request: StatusCode=404 -- Original Error: Code="ResourceGroupNotFound" Message="Resource group 'APIC-infra-mininet-fchazel-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]/csc-[ct_routerp_centralus_0_0]/nstoper	Operational State of the hcloud InstanceOper is down with [computeVirtualMachinesClientCreateOrUpdate: Failure sending request: StatusCode=404 -- Original Error: Code="ResourceGroupNotFound" Message="Resource group 'APIC-infra-mininet-fchazel-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]/csc-[ct_routerp_eastus_0_0]/license/oper	Operational State of the hcloud InstanceOper 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]/csc-[ct_routerp_centralus_1_0]/license/oper	Operational State of the hcloud InstanceOper is down with administrative-down	raised	Sep 11 2019 05:21:35pm
No	Major	F0101	topology/pool-1/node-1/sys/chp-[jdev/vrb]-[jdev/vrb]	Storage unit /dev/vrb 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 Network
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
canadaeast	unmanaged	N/A	N/A	N/A	N/A
brazilsouth	unmanaged	N/A	N/A	N/A	N/A
australiaeast	unmanaged	N/A	N/A	N/A	N/A
australiacentral2	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



## CHAPTER 6

# Deploying Layer 4 to Layer 7 Services

- [Overview, on page 95](#)
- [Example Use Cases, on page 103](#)
- [Guidelines and Limitations for Redirect, on page 116](#)
- [Adding a New CIDR to Overlay-2 Using the Cloud APIC GUI, on page 118](#)
- [Deploying a Service Graph, on page 119](#)

## Overview

The Cisco Cloud APIC enables you to deploy Layer 4 to Layer 7 service devices to the public cloud. The initial release (4.2(x)), supports Azure Application Gateway (Application Load Balancer) deployments in Azure. Beginning with release 5.0(2), Azure Load Balancer (Network Load Balancer) and Third Party Firewall deployments in Azure are supported.

Two types of Load Balancers are supported for Layer 4 - Layer 7 deployments in Azure:

- ALB refers to Azure Application gateway or Application Load balancer
- NLB refers to Azure Load balancer or Network Load balancer.

## About Service Graphs

A service graph is used to represent a set of Layer 4- Layer 7 service devices inserted between two or more pair of EPGs. EPGs can represent your applications running within a cloud (e.g. Cloud EPG) or internet (cloudExtEPG) or from other sites (e.g. on-prem or remote cloud sites). Layer 4- Layer 7 devices can be NLB, ALB or a cluster of Third party firewalls.

A service graph in conjunction with contracts (and filters) is used to specify communication between two EPGs. A cloud APIC automatically derives security rules (network security group/NSG and ASG) and forwarding routes (UDRs) based on the policy specified in Contract and Service Graph

Multiple service graphs can be specified to represent a different represent different traffic flows or topologies.

Following combinations are possible with service graphs:

- Same device can be used in multiple service graphs.
- Same service graph can be used between multiple consumer and provider EPGs.

By using a service graph, the user can specify the policy once and deploy the service chain within regions or inter-regions. Each time the graph is deployed, Cisco ACI takes care of changing the network configuration to enable the forwarding in the new logical topology.

For Third party firewalls, the configuration inside the device is not managed by cloud APIC.

A service graph represents the network using the following elements:

- **Service Graph Nodes**—A 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.
- **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.

## About Application Load Balancers

Application Load Balancer (also called Azure Application Gateway or ALB) is a Layer 7 load balancer, which balances the web traffic based on attributes like HTTP request, URL filtering etc. For more details please refer to [Microsoft Documentation](#).

In Cisco ACI, there are two ways to deploy an Application Load Balancer:

- **Internet-facing:** inserts the Application Load Balancer as a service between the consumer external EPG and the provider cloud EPG.
- **Internal-facing:** inserts the Application Load Balancer as a service between the consumer cloud EPG and the provider cloud EPG.

You can consume an Application Load Balancer using a service graph. A typical configuration involves:

- Creation of L4L7 device as Application Load Balancer
- Consume the ALB as a node in the service graph
- Creation of one or more listeners in EPG communication when a service graph is associated with a contract.

Listeners enable you to specify the ports and protocols (HTTP or HTTPS) that the Application Load Balancer 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.

An Application load balancer (ALB) should be in a separate subnet which should not be used to deploy other applications. Cloud APIC creates and attaches ALB's NSG to the subnet associated with the ALB. Cloud APIC supports Standard and Standard\_v2 SKUs of Azure Application Gateway.

## About Network Load Balancer

A Network Load Balancer (Azure Load Balancer or NLB) is a Layer 4 device that distributes the in-bound flow packets based on L4 ports. For more details, please refer to [Microsoft Documentation](#).

Similar to ALB, NLB can be deployed using a service graph. You can specify these actions by configuring one or more listeners.

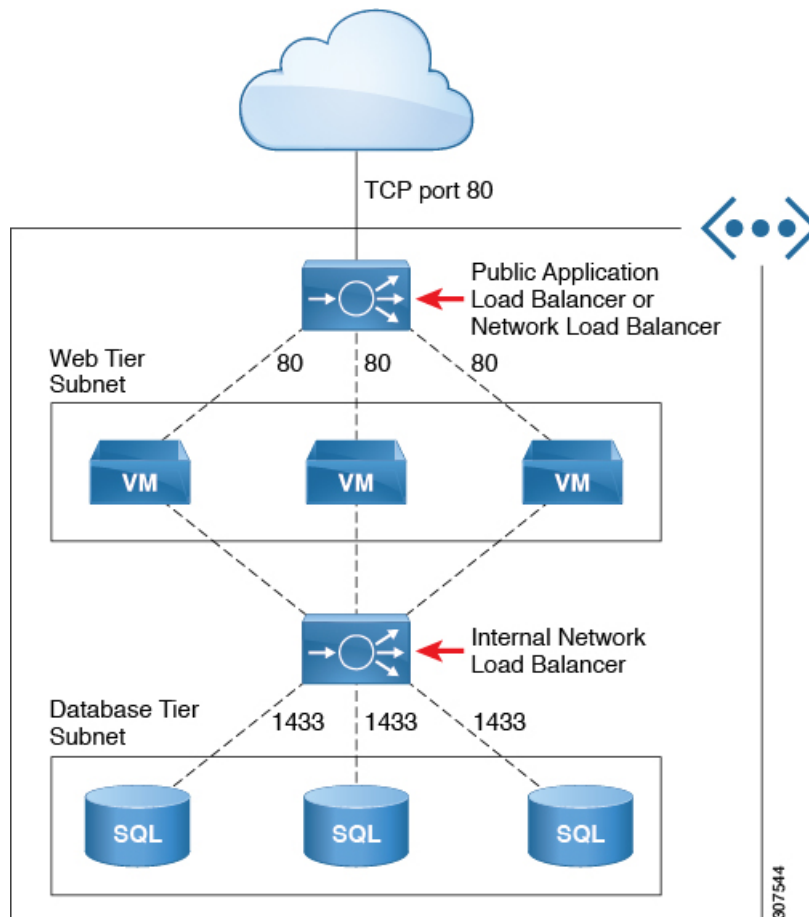
Listeners enable you to specify the ports and protocols (TCP or UDP) that the load balancer accepts and forwards traffic on. All listeners require you to configure at least one rule (a default rule, which does not have a condition). Rules enable you to specify the action that the load balancer takes when a condition is met. Unlike application gateway, here a rule can only forward traffic to specific port of the backend pool. NLB should be in a separate subnet similar to ALB. There are two modes of operation in Network load balancer:

- Forward mode: Traffic is forwarded from a specific listener port to the specified backend port.
- HA Port mode: Network load balancer will load balance TCP and UDP flows on all the ports simultaneously.

Cloud APIC supports Standard SKU Network Load Balancer only.

In Figure1, the frontend load balancer (ALB/NLB) - VM or firewall - backend load (ALB/NLB) balancer as a service are inserted between the consumer external EPG and the provider cloud EPG.

Figure 19: Internet-Facing and Internal-Facing Deployment



## Dynamic Server Attachment to Server Pool

Servers in provider EPG are dynamically added to the target groups. In Azure, the target groups are referenced as the backend pool. Listeners and rule configuration that define the frontend and backend protocol and port number, and load balancing action are provided by the user. When configuring listener rule as part of service graph configuration, user can select provider EPG for a given rule. The endpoints from that EPG would be dynamically added to the target group of the load balancer. You do not need to specify the endpoints or FQDN for the targets.

## About Inter-VNet Services

Beginning with Release 5.0(2), support is available for the deployment and automation of the inter-VNet services. This is both for the East-West and North-South use cases within the cloud.


Note the following considerations for this support:

- VNet peering needs to be configured for hub-spoke topology. For more information, refer to [Configuring VNet Peering for Cloud APIC for Azure](#).

- For multi-node services with redirect: The service device has to be present in the infra VNet. Service devices such as ALB fronting the provider can be present in the provider VNet.
- For multi-node service without redirect: The service device can be in the provider VNet or spread across the hub VNet and the provider VNet.
- Inter-VNet traffic is supported with an Application load balancer or Network load balancer in the infra VNet and the provider in a non-infra VNet. The VNets should be peered together and the load balancer and the provider should be from the same region.

## About Multinodes

Beginning with release 5.0(2), Multinode service graph is supported. Multinodes enable multiple deployment scenarios with service graphs.



Service devices that can be deployed are Application Load Balancer, Network Load Balancer and Third Party Firewall.

Two types of nodes are admitted in a graph.

- Non-redirect: Traffic is destined to service devices (Load Balancers, Thirdparty firewalls with DNAT and SNAT, Network Load Balancer).
- Redirect: Service device is a passthrough device (Network Load Balancer or Firewall).

## About Layer 4 to Layer 7 Service Redirect

Beginning with Release 5.0(2), the Layer 4 to Layer 7 Service Redirect feature is available for Cisco Cloud APIC, similar to the policy-based redirect (PBR) feature available for Cisco APIC. The Layer 4 to Layer 7 Service Redirect feature is configured using the **Redirect** option in the Cisco Cloud APIC.



**Note** Throughout this section, the term "consumer-to-provider" is sometimes used as a blanket term to describe traffic going from point A to point B, where a redirect service device might be inserted between those two points. However, this does not mean that only consumer-to-provider traffic is supported for redirect; traffic might also be from provider-to-consumer, such as in the use case described in [Spoke to Spoke, on page 105](#).

With redirect, policies are used to redirect traffic through specific service devices, where service devices can be deployed as a Network Load Balancer or a third-party firewall. This traffic isn't necessarily destined for the service device as part of the standard consumer-to-provider configuration; rather, you would configure the consumer-to-provider traffic as you normally would, and you would then configure service graphs to redirect that consumer-to-provider traffic to a specific service device.

Support for redirect for Cisco Cloud APIC is only available in conjunction with the VNet peering feature, taking advantage of the hub-and-spoke topology used in VNet peering. For more information on the VNet peering feature, see the [Configuring VNet Peering for Cloud APIC for Azure](#) document.

### About the Overlay-1 and Overlay-2 VRFs

The overlay-1 and overlay-2 VRFs are automatically created in the infra tenant for Cloud APIC. In the Azure portal, CIDRs and subnets from the overlay-1 and overlay-2 VRFs are deployed in the Azure cloud on the overlay-1 VNet. The overlay-2 VRF is used to hold additional CIDRs. You shouldn't consider overlay-2 as a separate VNet.

The following sections provide more information on the overlay-1 and overlay-2 VRFs.

#### Requirement for Separate VRFs in the Infra Hub

Prior to Release 5.0(2), the infra hub VNet was used to achieve transit routing functionality for inter-spoke communications within the site through CSRs in the hub, and to send VxLAN packets for EPG communication across sites.

There are situations where you might want to deploy a certain number of EPGs configured with shared services and layer 4 to layer 7 service graphs in a common hub that can be shared across spokes. In some situations, you might have multiple hub networks deployed separately (for example, for production, pre-production, and core services). You might want to deploy all of these hub networks in the same infra hub VNet (in the same infra cloud context profile), along with the existing cloud CSRs.

Thus, for these kind of requirements, you might need to split the hub VNet into multiple VRFs for network segmentation while keeping the security intact.

#### About the Infra Hub Services VRF (Overlay-2 VRF in the Infra VNet)

Beginning with Release 5.0(2), the overlay-2 VRF is now created in the infra tenant implicitly during the Cisco Cloud APIC bringup. In order to keep the network segmentation intact between the infra subnets used by the cloud site (for CSRs and network load balancers) and the user subnets deployed for shared services, different VRFs are used for infra subnets and user-deployed subnets:

- **Overlay-1:** Used for infra CIDRs for the cloud infra, along with Cisco Cloud Services Routers (CSRs), the infra network load balancer, and the Cisco Cloud APIC
- **Overlay-2:** Used for user CIDRs to deploy shared services, along with layer 4 to layer 7 service devices in the infra VNet (the overlay-1 VNet in the Azure cloud)

All the user-created EPGs in the infra tenant can only be mapped to the overlay-2 VRF in the infra VNet. You can add additional CIDRs and subnets to the existing infra VNet (the existing infra cloud context profile). They are implicitly mapped to overlay-2 VRF in the infra VNet, and are deployed in the overlay-1 VNet in the Azure cloud.

Prior to Release 5.0(2), any given cloud context profile would be mapped to a cloud resource of a specific VNet. All the subnets and associated route tables of the VNet would be have a one-to-one mapping with a single VRF. Beginning with Release 5.0(2), the cloud context profile of the infra VNet can be mapped to multiple VRFs (the overlay-1 and overlay-2 VRFs in the infra VNet).



In the cloud, the subnet's route table is the most granular entity for achieving network isolation. So all system-created cloud subnets of the overlay-1 VRF and the user-created subnets of the overlay-2 VRF will be mapped to separate route tables in the cloud for achieving the network segmentation.



**Note** On Azure cloud, you cannot add or delete CIDRs in a VNet when it has active peering with other VNets. Therefore, when you need to add more CIDRs to the infra VNet, you need to first disable VNet peering in it, which removes all the VNet peerings associated with the infra VNet. After adding new CIDRs to the infra VNet, you need to enable VNet peering again in the infra VNet.

You do not have to disable VNet peering if you are adding a new subnet in an existing CIDR in the hub VNet.

See [Adding a New CIDR to Overlay-2 Using the Cloud APIC GUI, on page 118](#) for more information.

## Passthrough Rules

When redirect is enabled, the rules in the NSGs (Network Security Groups) attached to the service devices are updated to permit traffic from consumer to provider. These rules are called "passthrough rules". In general, the passthrough rule is to permit traffic from consumer IP to provider IP. If the destination IP is an application load balancer (ALB) VIP, the rule is to permit traffic from consumer IP to the ALB VIP.

## Redirect Programming

Redirect programming depends on the classification of the destination EPG (tag-based or subnet-based):

- For a subnet-based EPG, subnets of the destination EPGs are used to program redirects
- For a tag-based EPGs, CIDRs of the destination VNet are used to program redirects

As a result of this, the redirect affects traffic from other EPGs going to the same destination in the redirect, even if the EPG is not part of the service graph with the redirect. Traffic from EPGs that are not part of the redirect will also get redirected to the service device.

The following table describes how redirect is programmed in different scenarios.

Consumer	Provider	Redirect on Consumer VNet	Redirect on Provider VNet
Tag-based	Tag-based	Redirect for the provider are the CIDRs of the provider's VNet	Redirect for the consumer are the CIDRs of the consumer's VNet
Tag-based	Subnet-based	Redirect for the provider are the subnets of the provider	Redirect for the consumer are the CIDRs of the consumer's VNet
Subnet-based	Tag-based	Redirect for the provider are the CIDRs of the provider's VNet	Redirect for the consumer are the subnets of the consumer
Subnet-based	Subnet-based	Redirect for the provider are the subnets of the provider	Redirect for the consumer are the subnets of the consumer

## Redirect Policy

To support the Layer 4 to Layer 7 Service Redirect feature, a new redirect flag is now available for service device connectors. The following table provides information on the existing and new flags for the service device connectors.

ConnType	Description
<b>redir</b>	This value means the service node is in redirect mode for that connection. This value is only available or valid for third-party firewalls and Network Load Balancers.
<b>snat</b>	This value tells the service graph that the service node is performing source NAT on traffic. This value is only available or valid for the provider connector of third-party firewalls and only on the provider connector of a node.
<b>snat_dnat</b>	This value tells the service graph that the service node is performing both source NAT and destination NAT on traffic. This value is only available or valid for the provider connector of third-party firewalls and only on the provider connector of a node.
<b>none</b>	Default value.

## Workflow for Configuring Redirect

Following is the typical workflow for configuring redirect:

1. Create one or more service devices to use with the service graph:
  - Network load balancer (NLB)
  - Application load balancer (ALB)
  - Third-party firewall
2. Create a service graph and select the appropriate service devices for this particular service graph.
 

You will configure redirect at this point in the procedures:

  - a. Drag and drop a network load balancer, application load balancer, or firewall icon to the **Drop Device** area to select that service device for the service graph.
  - b. To enable the redirect feature, in the **Service Node** window that appears, check the box next to the **Redirect** option under the **Consumer Connector Type** and/or under the **Provider Connector Type** areas, depending on where you want to enable the redirect function.




---

**Note** Even though you might have an application load balancer in the service graph, you cannot enable redirect on an application load balancer service device.

---

- c. Complete the remaining configurations in the **Service Node** window, then click **Add**.
3. Configure the EPG communication, where you create a contract between the consumer and the provider EPGs.
4. Attach the service graph to the contract.
5. Configure the service device parameters.

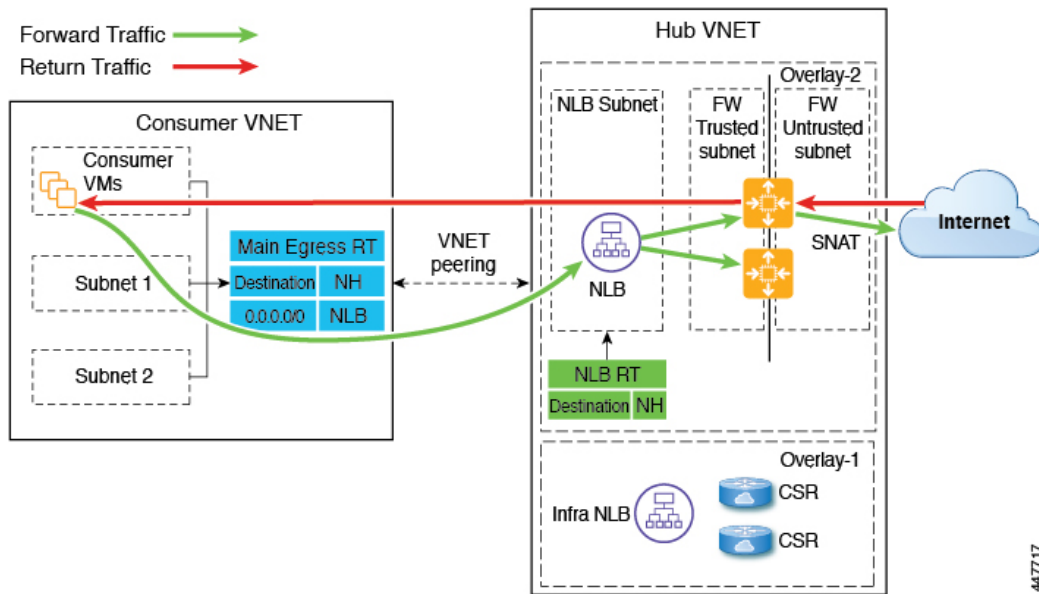
## Example Use Cases

Following are several example use cases:


- [Spoke to Internet, on page 103](#)
- [Spoke to Spoke, on page 105](#)
- [Inter-Region Spoke to Spoke, on page 108](#)
- [Internet to Spoke \(Inter-VRF\), on page 110](#)
- [Consumer and Provider EPGs in Two Separate VNets, on page 112](#)
- [Hub VNet with Consumer and Provider EPGs in Two Separate VNets, on page 114](#)

### Spoke to Internet


In this use case, the consumer VNet (with consumer VMs) and the hub VNet are peered using VNet peering. A network load balancer is also deployed, fronting two firewalls for scaling. In this use case, the consumer VMs need access to the internet for a certain reason, such as patch updates. In the consumer VNet, the route table is modified to include a redirect for the internet in this case, and traffic is redirected to the NLB in front of firewalls in the hub VNet. Any traffic from this consumer that is part of the service graph that is going to the internet goes to the NLB as the next-hop. With VNet peering, traffic first goes to the NLB, then the NLB forwards the traffic to one of the firewalls in the back end. The firewalls also perform source network address translation (SNAT) when sending traffic to the internet.



The following figure shows the packet flow for this use case.



The following figure shows the service graph for this use case.

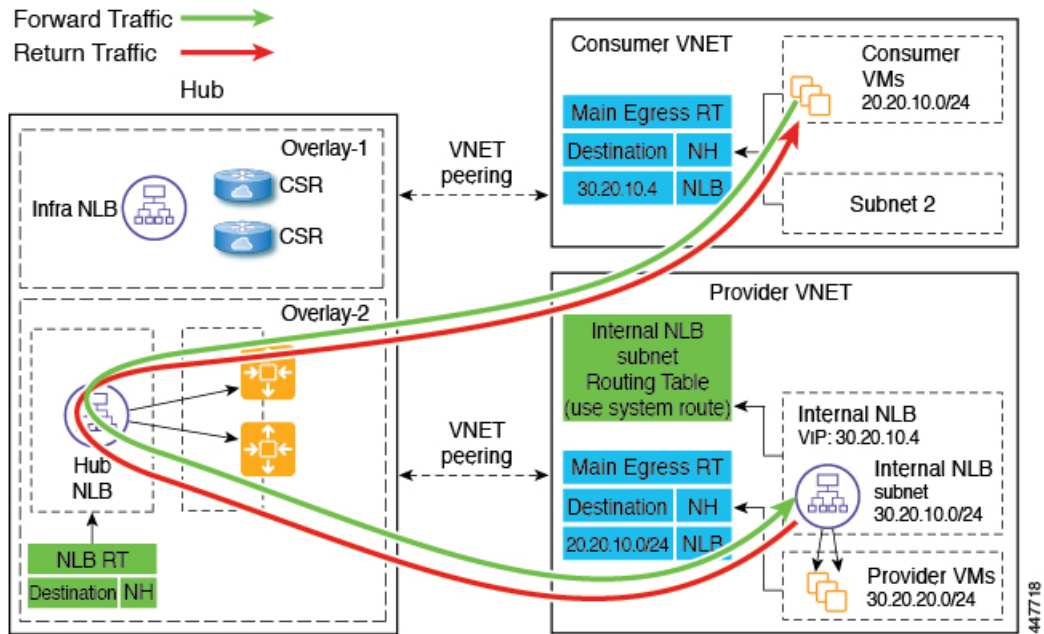


As part of the redirect configuration for this use case, you would make the following selections:


- In the **Create Device** window
  - In the **Tenant** field, choose the **infra** tenant.
  - Choose the type of service device in the **Service Type** field:
    - Choose **Network Load Balancer** as the **Service Type**, and in the **Subnets** area, click **Add Subnet**, then choose the appropriate region, cloud context profile, and the subnet created in the overlay-2 VRF.
    - Choose **Third-Party Firewall** as the **Service Type**, and in the **VRF** field, choose the **overlay-2** VRF.
  
- In the **Create Service Graph** window, drag-and-drop the following service devices, in this order:
  - Network Load Balancer
  - Third-Party Firewall
  
- In the **Service Node** window for the Network Load Balancer:
  - In the **Consumer Connector Type** field, place a check in the box next to the **Redirect** option to enable the redirect function on the consumer side.
  - In the **Provider Connector Type** field, leave the boxes unchecked.
  
- In the **Service Node** window for the Third-Party Firewall:
  - In the **Consumer Connector Type** field, leave the boxes unchecked.
  - Because the firewall performs SNAT when sending traffic to the internet in this use case, in the **Provider Connector Type** field, place a check in the box next to the **SNAT** option.

### Spoke to Spoke


In this use case, traffic flows from spoke to spoke, through the hub firewall fronted by a hub NLB. Consumer endpoints are in the consumer VNet, and the provider VNet has VMs fronted by an internal NLB. The egress route table is modified in the consumer and provider VNets so that traffic is redirected to the firewall device fronted by the NLB. Redirect is applied in both directions in this use case. The NLB must have a dedicated subnet in this case.



The following figure shows the packet flow for this use case.



The following figure shows the service graph for this use case.

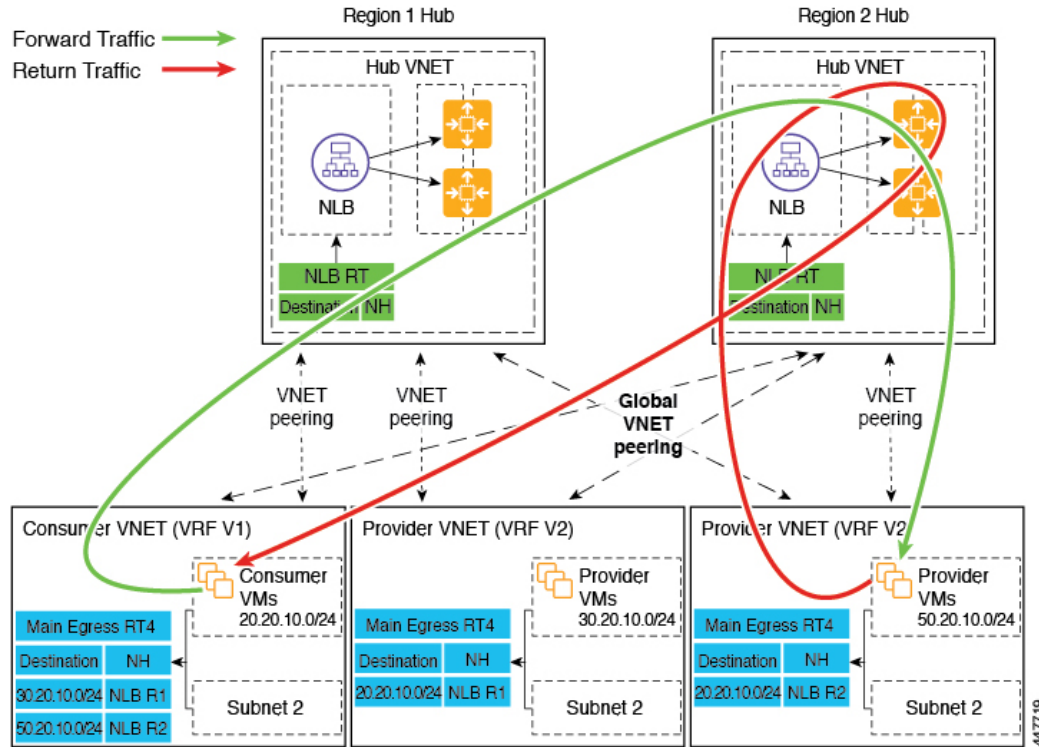


As part of the redirect configuration for this use case, you would make the following selections:

- In the **Create Device** window, first create the service devices for the hub VNet:
  - In the **Tenant** field, choose the **infra** tenant.
  - Choose the type of service device in the **Service Type** field:
    - Choose **Network Load Balancer** as the **Service Type**, and in the **Subnets** area, click **Add Subnet**, then choose the appropriate region, cloud context profile, and the subnet created in the overlay-2 VRF.
    - Choose **Third-Party Firewall** as the **Service Type**, and in the **VRF** field, choose the **overlay-2** VRF.
- In the **Create Device** window, next create the service devices for the provider VNet:
  - In the **Tenant** field, choose the provider tenant.
  - In the **Service Type** field, choose **Network Load Balancer**, and in the **Subnets** area, click **Add Subnet**, then choose the appropriate region, cloud context profile, and the subnet for the provider VRF.
- In the **Create Service Graph** window, drag-and-drop the following service devices, in this order:
  - Network Load Balancer (for the hub VNet)
  - Third-Party Firewall (for the hub VNet)
  - Network Load Balancer (for the provider VNet)
- In the **Service Node** window for the Network Load Balancer in the hub VNet:
  - In the **Consumer Connector Type** field, place a check in the box next to the **Redirect** option to enable the redirect function on the consumer side.
  - In the **Provider Connector Type** field, place a check in the box next to the **Redirect** option to enable the redirect function on the provider side.
- In the **Service Node** window for the Third-Party Firewall, leave the boxes unchecked for the **Consumer Connector Type** and the **Provider Connector Type**.
- In the **Service Node** window for the Network Load Balancer in the provider VNet, leave the boxes unchecked for the **Consumer Connector Type** and the **Provider Connector Type**.

## Inter-Region Spoke to Spoke

In this use case, both regions must have service devices. The consumer VNet is in region 1, the provider is stretched across both regions (regions 1 and 2), and some endpoints are in region 1 and some endpoints are in region 2. Different redirects are programmed for local provider endpoints and for remote region endpoints. In this case, the firewall that is used will be the firewall that is closest to the provider endpoint side.



For example, consider the two subnets in the consumer VNet (VRF 1) egress route table (RT):

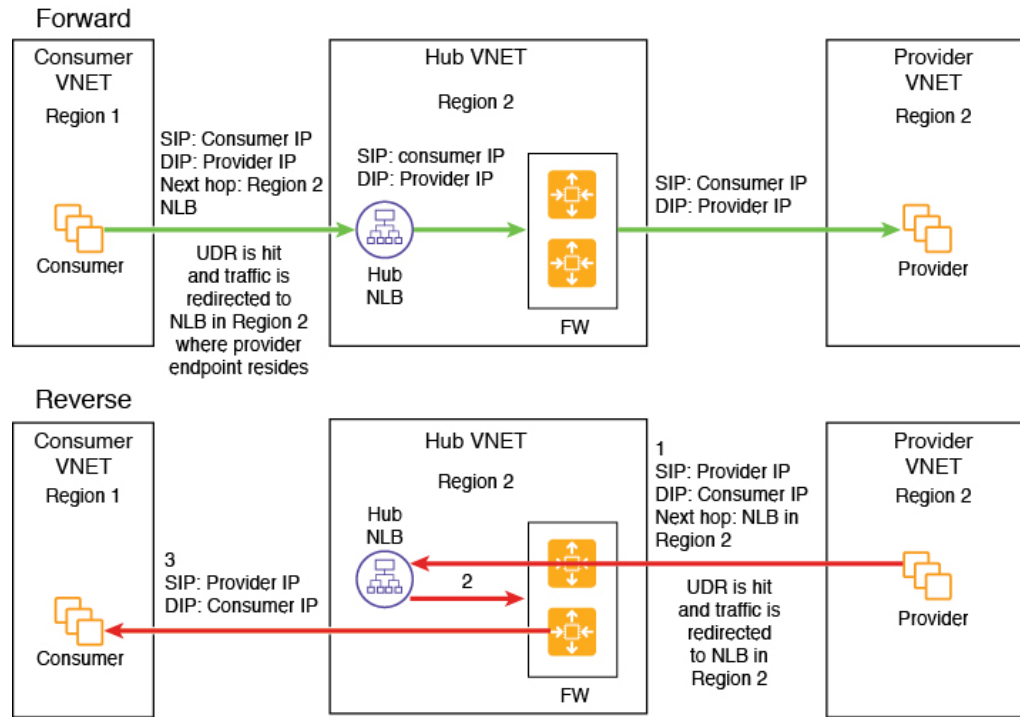
- 30.20.10.0/24 (NLB in region 1 [R1])
- 50.20.10.0/24 (NLB in region 2 [R2])

Assume the consumer wants to send traffic to the provider VMs 30.20.10.0/24, which are local to it. In that case, traffic will get redirected to the region 1 hub NLB and firewall, and will then go to the provider.


Now assume the consumer wants to send traffic to the provider VMs 50.20.10.0/24. In this case, the traffic will get redirected to the region 2 hub NLB and firewall, because that firewall is local to the provider endpoint.

The following figure shows the packet flow for this use case.





The following figure shows the service graph for this use case.



As part of the redirect configuration for this use case, you would make the following selections:

- In the **Create Device** window, first create the service devices for the hub VNet:
  - In the **Tenant** field, choose the **infra** tenant.
  - Choose the type of service device in the **Service Type** field:
    - Choose **Network Load Balancer** as the **Service Type**, and in the **Subnets** area, click **Add Subnet**, then choose the appropriate region, cloud context profile, and the subnet created in the overlay-2 VRF.
    - Choose **Third-Party Firewall** as the **Service Type**, and in the **VRF** field, choose the **overlay-2** VRF.
- In the **Create Service Graph** window, drag-and-drop the following service devices, in this order:
  - Network Load Balancer
  - Third-Party Firewall

- In the **Service Node** window for the hub NLB:
  - In the **Consumer Connector Type** field, place a check in the box next to the **Redirect** option to enable the redirect function on the consumer side.
  - In the **Provider Connector Type** field, place a check in the box next to the **Redirect** option to enable the redirect function on the provider side.
- In the **Service Node** window for the Third-Party Firewall, leave the boxes unchecked for the **Consumer Connector Type** and the **Provider Connector Type**.

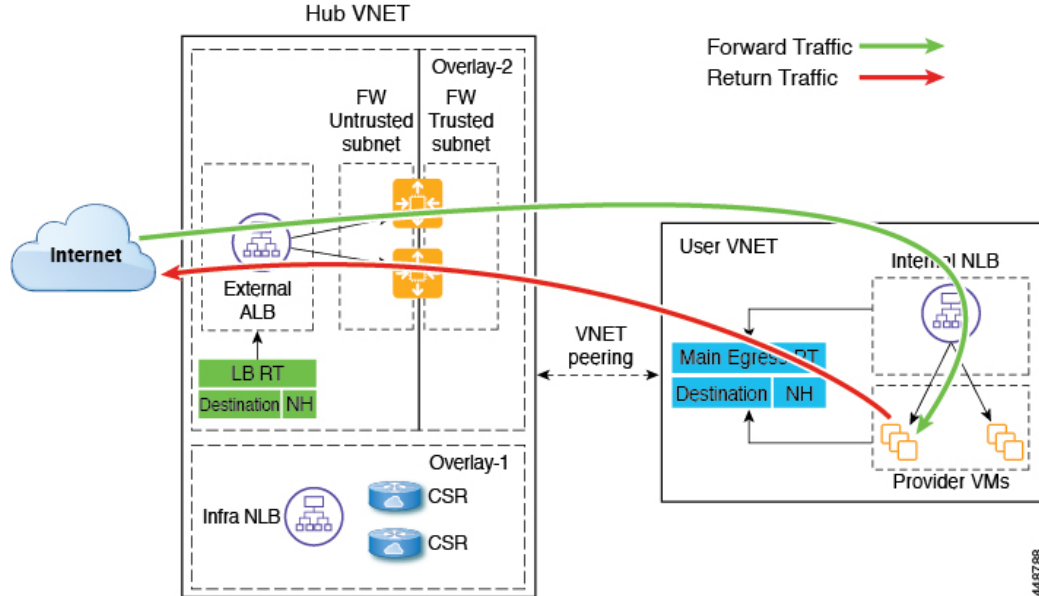
### Internet to Spoke (Inter-VRF)

In this use case, traffic coming from the internet needs to go through the firewall before hitting the provider endpoints. Redirect is not used in this use case.

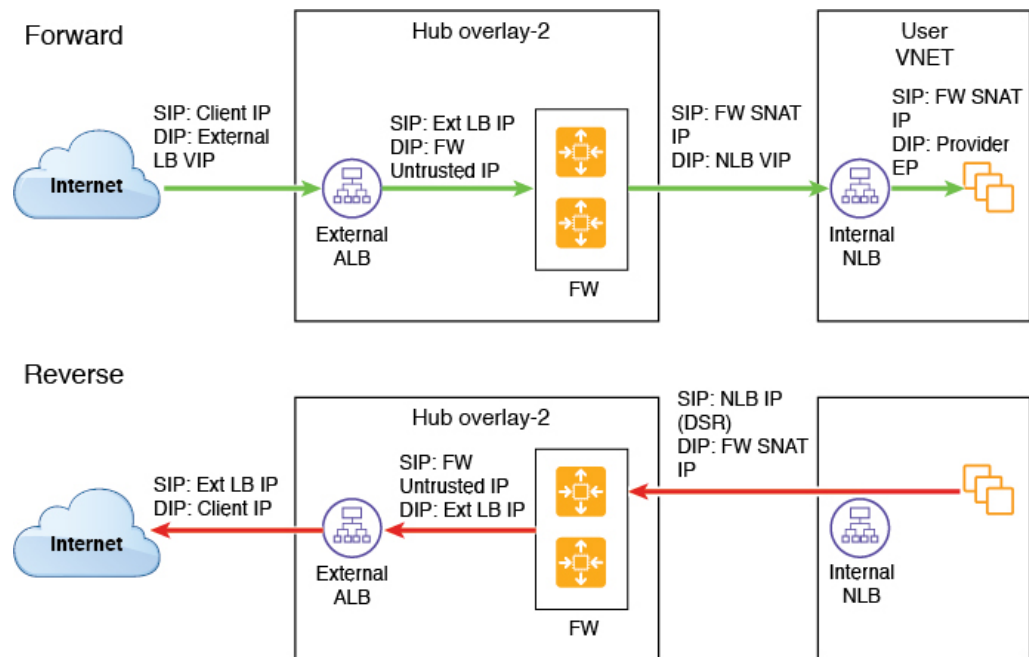


**Note** The general term "external load balancer" is used in this section because in this use case, the external load balancer could be either an NLB or an ALB. The following examples provide configurations using an ALB, but keep in mind that the external load balancer could be an NLB instead.


The external load balancer exposes the service through VIP. Internet traffic is directed to that VIP, then external load balancers direct traffic to the firewalls in the backend pool (the external load balancers have the firewall's untrusted interface as its backend pool). The firewall performs SNAT and DNAT, and the traffic goes to the internal NLB VIP. The internal NLB then sends traffic to one of the provider endpoints.



The following figure shows the packet flow for this use case.



The following figure shows the service graph for this use case.



As part of the redirect configuration for this use case, you would make the following selections:

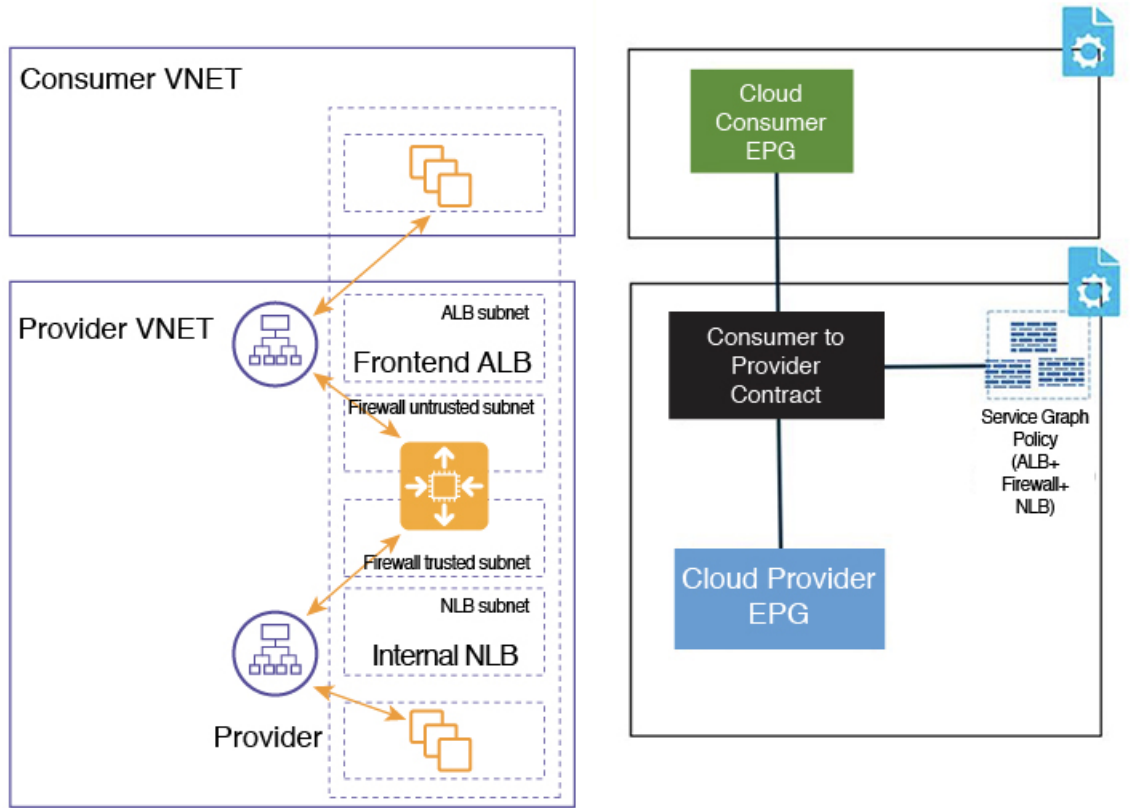
- In the **Create Device** window, first create the service devices for the hub VNet:
  - In the **Tenant** field, choose the **infra** tenant.
  - Choose the type of service device in the **Service Type** field:
    - Choose **Application Load Balancer** or **Network Load Balancer** as the **Service Type**, and in the **Subnets** area, click **Add Subnet**, then choose the appropriate region, cloud context profile, and the subnet created in the overlay-2 VRF.
    - Choose **Third-Party Firewall** as the **Service Type**, and in the **VRF** field, choose the **overlay-2** VRF.
- In the **Create Device** window, next create the service devices for the provider VNet:
  - In the **Tenant** field, choose the provider tenant.
  - In the **Service Type** field, choose **Network Load Balancer**, and in the **Subnets** area, click **Add Subnet**, then choose the appropriate region, cloud context profile, and the subnet for the provider VRF.

- In the **Create Service Graph** window, drag-and-drop the following service devices, in this order:
  - Network Load Balancer or Application Load Balancer (for the hub VNet)
  - Third-Party Firewall (for the hub VNet)
  - Network Load Balancer (for the provider VNet)
- In the **Service Node** window for the Network Load Balancer or Application Load Balancer for the hub VNet, leave the boxes unchecked for the **Consumer Connector Type** and the **Provider Connector Type**.
- In the **Service Node** window for the Third-Party Firewall:
  - In the **Consumer Connector Type** field, leave the boxes unchecked.
  - Because the firewall performs SNAT and DNAT when sending traffic to the internet in this use case, in the **Provider Connector Type** field, place a check in the box next to the **SNAT** and **DNAT** options.
- In the **Service Node** window for the Network Load Balancer for the provider VNet, leave the boxes unchecked for the **Consumer Connector Type** and the **Provider Connector Type**.

### Consumer and Provider EPGs in Two Separate VNets

This use case is an example configuration with two VNets, with a consumer EPG and provider EPG in separate VNets.

- A frontend ALB, firewall, and internal NLB are inserted between the consumer and provider EPGs.
- A consumer endpoint sends traffic to the frontend ALB VIP and it is forwarded to the firewall.
- The firewall performs SNAT and DNAT, and the traffic flows to internal NLB VIP.
- The internal NLB load balances the traffic to the backend provider endpoints.




In the figure:

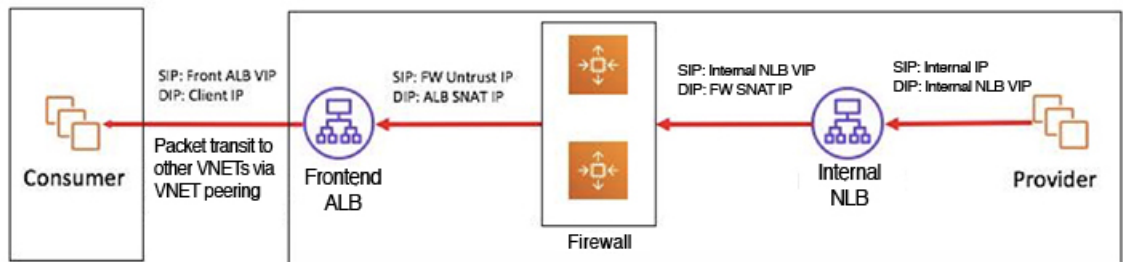
- The consumer EPG is in a consumer VNet.
- The provider EPG and all the service devices are in the provider VNet.
- The application load balancer, network load balancer, and firewall need to have their own subnet in the VNet.

Packet flow for both the directions is shown in the following figure:

## Forward



## Reverse

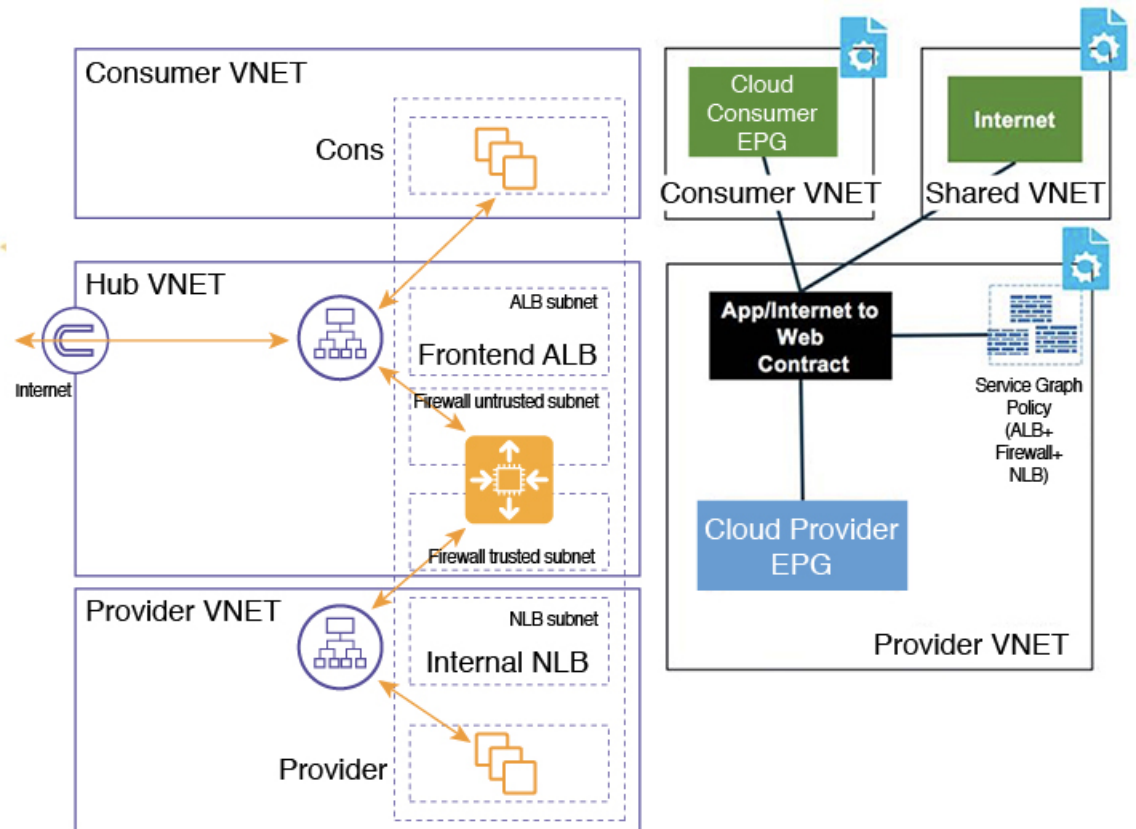


503031

**Hub VNet with Consumer and Provider EPGs in Two Separate VNets**

This use case is an example configuration with three VNets: a hub VNet, and a consumer EPG and provider EPG in two separate VNets.

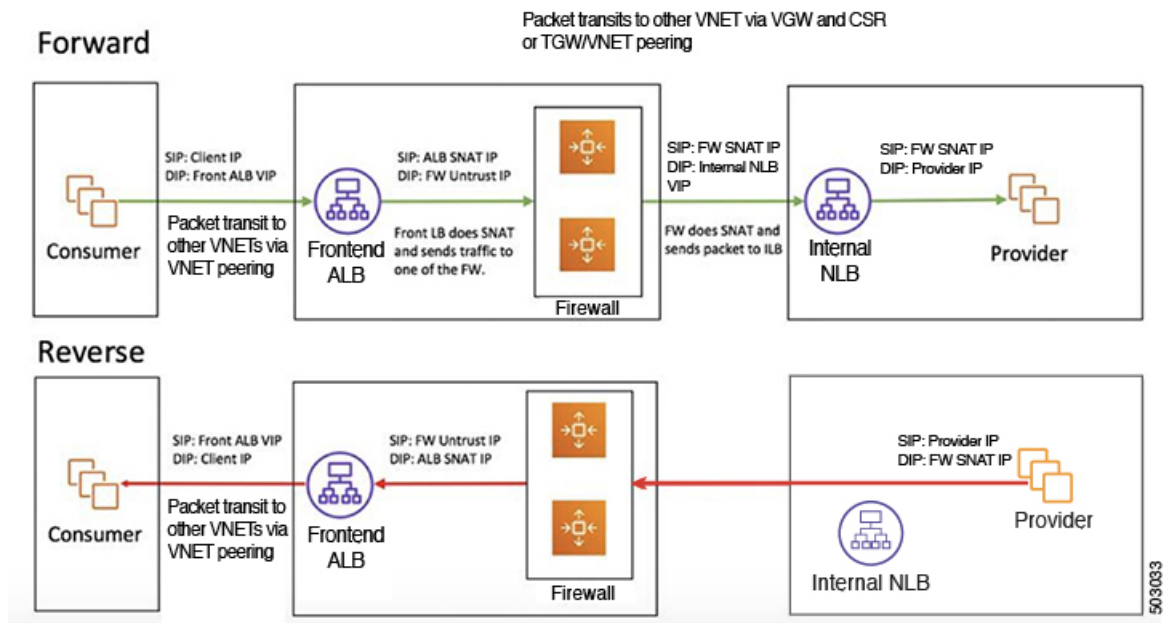
- A frontend ALB and firewall are inserted within the hub VNet, which is between the consumer and provider EPGs.
- An internal NLB is inserted in the provider EPG.
- A consumer endpoint sends traffic to the frontend ALB VIP and it is forwarded to the firewall.
- The firewall performs SNAT and DNAT, and the traffic flows to internal NLB VIP.
- The internal NLB load balances the traffic to the backend provider endpoints.



In the figure:

- The consumer EPG is in a consumer VNet.
- The provider EPG and the internal NLB are in the provider VNet.
- The frontend ALB and firewall are in the hub VNet
- The application load balancer, network load balancer, and firewall need to have their own subnet in the VNet.

Packet flow for both the direction is shown in the following figure:



## Guidelines and Limitations for Redirect

Following are the guidelines and limitations for redirect:

- All the Layer 4 - Layer 7 service devices should have their own dedicated subnet.
- Intra VRF Layer 4 - Layer 7 redirection within a region:
  - Layer 4 - Layer 7 redirect is not supported for east-west deployment when the consumer EPG and provider EPG are in the same VNet.
  - Layer 4 - Layer 7 redirect is supported for north-south deployment if the external EPG is a provider EPG, regardless of whether the consumer EPG and provider EPG are in same VNet or not.
- Intra-VRF Layer 4 - Layer 7 redirection across regions:
  - Inter-Region Layer 4 - Layer 7 redirection are supported. However, the Consumer EPG and the Provider EPG should not stretch.
  - A region shouldn't have both a consumer EPG and a provider EPG in the same VRF. For example, if region 1 has a consumer EPG only and region 2 has a provider EPG only, this is supported, but region 1 can't have both the consumer EPG and the provider EPG.
  - Consumer and Provider EPG should be a subnet-based EPG.
- For the inter-region service graphs with Layer 4 - Layer 7 redirection, service devices should be deployed in the provider EPG's region. If provider EPG is stretched across regions, service devices should be deployed in each region .
- For the external EPG as provider, service devices need to be deployed in the region local to consumer EPG. If the consumer EPG is stretched across regions, service devices should be deployed in each region.



- Between a consumer VNet and a provider EPG, only one redirect device can be inserted through a service graph. For example, if consumer EPG1 and consumer EPG2 are in a consumer VNet, and a provider EPG3 is in a provider VNet, you must use the same redirect device for a contract between EPG1 and EPG3, and a contract between EPG2 and EPG3.



**Note** The limitation is because of the cloud provider allows only one next hop for a given destination in user-defined routes.

- The following table provides information on the specific redirect configurations that are supported or unsupported, where:
  - NLB stands for network load balancer
  - ALB stands for application load balancer
  - FW stands for firewall

Service Chain Option	Spoke-to-Spoke		Spoke-to-External (consumer is spoke)		External-to-Spoke (consumer is external)	
	Intra-VNet	Inter-VNet	Intra-VNet	Inter-VNet	Intra-VNet	Inter-VNet
NLB/ALB <sup>1</sup>	Supported	Supported	Not supported	Not supported	Supported	Supported
FW (no SNAT) <sup>2</sup>	Not supported	Supported	Not supported	Not supported	Not supported	Not supported
FW (w/SNAT) <sup>3</sup>	Supported	Supported	Supported	Supported	Not supported	Not supported
<ul style="list-style-type: none"> <li>• NLB<sup>2</sup>-FW(no SNAT)<sup>1</sup></li> <li>• NLB<sup>2</sup>-FW(no SNAT)<sup>1</sup>-NLB/ALB<sup>1</sup></li> </ul>	Not supported	Supported	Not supported	Not supported	Not supported	Not supported
NLB <sup>4</sup> -FW(SNAT) <sup>5</sup>	Not supported	Supported	Supported	Supported	Not supported	Not supported
NLB/ALB <sup>1</sup> -FW(SNAT+DNAT) <sup>6</sup> -NLB/ALB <sup>1</sup> (No redirection)	Supported	Supported	Supported	Supported	Supported	Supported

<sup>1</sup> Unchecked on both consumer and provider connector or options are not applicable for ALB.

<sup>2</sup> Redirect is enabled on both consumer and provider connector.

<sup>3</sup> Redirect is enabled on consumer connector. SNAT is enabled on provider connector.

<sup>4</sup> Redirect is enabled on consumer connector. Unchecked on provider connector.

<sup>5</sup> Unchecked on consumer connector. SNAT is enabled on provider connector.

<sup>6</sup> Unchecked on consumer connector. SNAT+DNAT is enabled on provider connector.

## Adding a New CIDR to Overlay-2 Using the Cloud APIC GUI

After an installation, you will see overlay-1 and overlay-2 in the Cisco Cloud APIC. However, on the Azure portal, you will only see overlay-1. This is because overlay-2 is simply a logical extension of overlay-1, and is used to hold additional the CIDRs that you might need if you are deploying firewalls or load balancers on the infra VNet. This section provides instructions for adding new CIDRs to overlay-2.

In some situations, you might have to disable VNet peering before adding new CIDRs or editing existing CIDRs in overlay-2. This is due to a limitation in Azure, where you cannot update a CIDR on a VNet if it has active VNet peerings. To add the CIDRs, you first have to remove VNet peerings for that VNet, then you can update the CIDRs. Once you have updated the CIDRs, you can then re-enable the VNet peerings.

These procedures provide instructions for disabling Hub Network Peering, which removes all of the VNet peerings associated with a particular infra VNet.

- If you have an additional CIDR already created on the infra VNet, but you simply need to add additional subnets to that existing CIDR, you do not have to disable Hub Network Peering for that particular infra VNet before adding those subnets. To add additional subnets to an existing CIDR:
  1. Navigate to the appropriate cloud context profile in that case (**Application Management > Cloud Context Profiles**).
  2. Double-click the cloud context profile where you want to add a subnet to an existing CIDR, then go to [Step 10, on page 119](#) to add the new subnets to an existing CIDR.
- If you are adding a new CIDR in the infra VNet, or if you are deleting a CIDR or editing a CIDR in the infra VNet in some other way (other than adding subnets), then you must disable Hub Network Peering for that particular infra VNet. You will then re-enable Hub Network Peering again after you have added the CIDR. The following procedure provides those instructions.

- 
- Step 1** Log in to the Cloud APIC, if you are not logged in already.
- Step 2** In the left navigation bar, navigate to **Application Management > Cloud Context Profiles**.  
The existing cloud context profiles are displayed.
- Step 3** Double-click the cloud context profile where you want to disable Hub Network Peering.  
The overview window for that cloud context profile appears. You should see **Enabled** in the **Hub Network Peering** area in this overview window, which indicates that Hub Network Peering is enabled.
- Step 4** Click the pencil icon to edit this cloud context profile.  
The **Edit Cloud Context Profile** window appears.
- Step 5** In the **Edit Cloud Context Profile** window, locate the **Hub Network Peering** field and click the check box to remove the checkmark from the **Enabled** field.  
Disabling the **Hub Network Peering** option does not remove VNet peering at the global level, but rather removes all of the VNet peerings associated with this particular infra VNet.
- Step 6** Click **Save**.  
The overview window for that cloud context profile appears again. You should see **Disabled** in the **Hub Network Peering** area in this overview window, which indicates that Hub Network Peering is now disabled.

- Step 7** To add a new CIDR, click the pencil icon to edit this cloud context profile again.  
The **Edit Cloud Context Profile** window appears again.
- Step 8** Click **Add CIDR**.  
The **Add CIDR** dialog box appears.
- Step 9** Add the new CIDR in the **CIDR Block Range** field.  
Do not click the box in the **Primary** field (do not put a check in the box next to **yes** in the **Primary** field).
- Step 10** Click **Add Subnet** and enter the necessary subnet addresses in the **Address** field.  
Continue to click **Add Subnet** for additional subnets, if necessary.
- Step 11** When you have finished adding all of the necessary information in the **Add CIDR** window, click **Add**.  
The **Edit Cloud Context Profile** window appears again.
- Step 12** Confirm the information in the **Edit Cloud Context Profile** window, then click **Save**.  
The overview window for that cloud context profile appears. You should now see the new CIDR listed in the **CIDR Block Range** area.
- Step 13** If you disabled Hub Network Peering at the beginning of these procedures, re-enable it at this time.
- Click the pencil icon to edit this cloud context profile.  
The **Edit Cloud Context Profile** window appears.
  - In the **Edit Cloud Context Profile** window, locate the **Hub Network Peering** field and click the check box to add the checkmark in the **Enabled** field to re-enable VNet peerings for this particular infra VNet.
  - Click **Save**.  
The overview window for that cloud context profile appears again. You should see **Enabled** in the **Hub Network Peering** area in this overview window, which indicates that Hub Network Peering is now re-enabled again.
- As described previously, if you were to go to the Azure portal at this point, you will see any additional CIDRs and subnets that you added in these procedures in the overlay-1 VNet in Azure, which is the correct and expected behavior.
- 

## 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.

The Service graph can be deployed in two ways:

- Single node service graph: Only one device is deployed.
- Multinode service graph: Upto three nodes can be added to the service chain.

Before you can deploy a service graph in either a single node or multinode, you must configure the following:

1. A tenant
2. An application profile

3. A consumer EPG
4. A provider EPG
5. A cloud context profile
6. A contract with a filter

## Deploying a Service Graph Using the GUI

The following sections describe how to deploy a service graph using the GUI.

### Creating Service Devices Using The Cloud APIC GUI

#### Before you begin

This section explains how to create service devices that can be used in a service graph through 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 **Services > Devices > Create Device**. The **Create Device** dialog box appears.

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

**Table 29: Create Device Dialog Box Fields for Application Load Balancer**

Properties	Description
<b>General</b>	
<b>Name</b>	Enter the name of the device.
<b>Tenant</b>	To choose a tenant: <ol style="list-style-type: none"> <li>a. Click <b>Select Tenant</b>. The <b>Select Tenant</b> dialog appears.</li> <li>b. From the column on the left, click to choose a tenant.</li> <li>c. Click <b>Select</b>. You return to the <b>Create Device</b> dialog box.</li> </ol>
<b>Settings</b>	
<b>Service Type</b>	Choose the device type: <ul style="list-style-type: none"> <li>• Application Load Balancer</li> </ul>

Properties	Description
<b>ALB SKU</b>	Choose from: <ul style="list-style-type: none"> <li>• Standard</li> <li>• Standard V2</li> </ul>
<b>VM Instance Count</b>	Enter a number in the <i>VM Instance Count</i> text box. <b>Note</b> This is applicable only for the Application Gateway.
<b>VM Instance Size</b>	Click the radio button for the VM instance size you want to choose: <b>large</b> , <b>medium</b> , or <b>small</b> . <b>Note</b> This is applicable only for the Application Gateway.
<b>Scheme</b>	Choose <b>Internet Facing</b> or <b>Internal</b> . <ul style="list-style-type: none"> <li>• <b>Internet Facing</b>— This is used for configuring a public IP for the balancer. This is assigned by Azure.</li> <li>• <b>Internal</b>—Click to choose either <b>Dynamic</b> or <b>Static</b> under IP Address Assignment. <ul style="list-style-type: none"> <li>• <b>Dynamic</b>—Dynamic IP addresses are assigned by Azure. Dynamic IP addresses change each time the VMs boot up.</li> <li>• <b>Static</b>—Enter an IP address based on the CIDRs defined in Cloud Context Profile and check that the IP address is in the same subnet as the ALB.  ALB SKU Standard supports static and dynamic IP addresses. ALB SKU Standard V2 support static IP addresses only.</li> </ul> </li> </ul>
<b>Subnet</b>	To choose a subnet: <ol style="list-style-type: none"> <li>Click <b>Select Region</b>. The <b>Select Region</b> dialog box appears. From the <b>Select Region</b> dialog, click to choose a region in the left column then click <b>Select</b>.</li> <li>Click <b>Select Cloud Context Profile</b>. The <b>Select Cloud Context Profile</b> dialog box appears.</li> <li>Click <b>Select Subnet</b>. The <b>Select Subnet</b> dialog box appears. The Static IP Addresses text box is displayed. Enter the IP address of the load balancer. Click the tick mark on the right to confirm.</li> <li>To add additional subnets, repeat steps a-c.</li> </ol>

**Step 5** Click **Save** when finished.

**Step 6** The **Create Service Graph** dialog box appears. Click on the **Create another Application Load Balancer** to create another device. The **Create Device** dialog box appears.

**Note** The UI usually asks to create a previously created device. However, on clicking it we return back to the **Create Device** page. Here we can choose the device that needs to be created. The first device should never be the Third Party Firewall.

**Step 7** Enter the appropriate values in each field as listed in the following *Create Device Dialog Box Fields* table then continue.

**Table 30: Create Device Dialog Box Fields for Third party firewall**

Properties	Description
<b>General</b>	
<b>Name</b>	Enter the name of the device.
<b>Settings</b>	
<b>Service Type</b>	Choose the device type: <ul style="list-style-type: none"> <li>• Third party firewall</li> </ul> <p><b>Note</b> Third party firewall cannot be the first device in a multinode service graph.</p>
<b>VRF</b>	To choose a VRF: <ol style="list-style-type: none"> <li>Click <b>Select VRF</b>. The <b>Select VRF</b> dialog box appears.</li> <li>From the <b>Select VRF</b> dialog, click to choose a VRF in the left column then click <b>Select</b>.</li> </ol>
<b>Interfaces</b>	Click <b>Add Interface selectors</b> <ol style="list-style-type: none"> <li>In the <b>Settings</b> page, enter the name of the interface.</li> <li>Click <b>Add Interface</b>.</li> <li>Enter the name of the interface selector.</li> <li>Click on <b>Match Expressions</b> and select <ul style="list-style-type: none"> <li>• the <b>Key</b>: This can be IP, region or a custom based tag selector.</li> <li>• <b>Operator</b>: This can be equal, not equals, in, not in, has key, or does not have key.</li> <li>• <b>Value</b>: IP address of the app, web, internal network, management network, or external network.</li> </ul> </li> <li>Click <b>Add</b>.</li> <li>Repeat steps a - d to add more interfaces.</li> </ol>

**Step 8** Click **Save** when finished.

**Step 9** The **Create Service Graph** dialog box appears. Click on the **Create another Third Party Firewall** to create another device. The **Create Device** dialog box appears.

**Note** The UI usually asks to create a previously created device. However, on clicking it we return back to the **Create Device** page. Here we can choose the device that needs to be created. The first device should never be the Third Party Firewall.

**Step 10** Enter the appropriate values in each field as listed in the following *Create Device Dialog Box Fields* table then continue.

Table 31: Create Device Dialog Box Fields for Network Load Balancer

Properties	Description
<b>General</b>	
<b>Name</b>	Enter the name of the load balancer.
<b>Settings</b>	
<b>Service Type</b>	Choose the device type: <ul style="list-style-type: none"> <li>• Network Load Balancer</li> </ul>
If you are choosing Network Load Balancer, use the steps below.	
<b>Scheme</b>	Choose <b>Internet Facing</b> or <b>Internal</b> . <ul style="list-style-type: none"> <li>• <b>Internet Facing</b>— This is used for configuring a public IP for the balancer. This is assigned by Azure.</li> <li>• <b>Internal</b>—Click to choose either <b>Dynamic</b> or <b>Static</b> under IP Address Assignment. <ul style="list-style-type: none"> <li>• <b>Dynamic</b>—Dynamic IP addresses are assigned by Azure. Dynamic IP addresses change each time the VMs boot up.</li> <li>• <b>Static</b>—Enter an IP address based on the CIDRs defined in Cloud Context Profile and check that the IP address is in the same subnet as the NLB. Static IP addresses are associated to load balancers.</li> </ul> </li> </ul> <p><b>Note</b> Cloud APIC creates standard SKU NLBs only.</p>
<b>Subnet</b>	To choose a subnet: <ol style="list-style-type: none"> <li>Click <b>Select Region</b>. The <b>Select Region</b> dialog box appears. From the <b>Select Region</b> dialog, click to choose a region in the left column then click <b>Select</b>.</li> <li>Click <b>Select Cloud Context Profile</b>. The <b>Select Cloud Context Profile</b> dialog box appears.</li> <li>Click <b>Select Subnet</b>. The <b>Select Subnet</b> dialog box appears. The Static IP Addresses text box is displayed. Enter the IP address of the load balancer. Click the tick mark on the right to confirm.</li> <li>To add additional subnets, repeat steps a-c.</li> </ol>

**Step 11** 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 for a single node or a multinode, using the Cisco Cloud APIC GUI .

**Before you begin**

You have already created the devices.

- 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 **Services > Service Graph > Create Service Graph**. The **Create Service Graph** pop-up appears. Click on **Let's Get Started**.
- Step 4** Enter the appropriate values in each field as listed in the following *Create Service Graph Dialog Box Fields* table then continue.

**Table 32: Create Service Graph Dialog Box Fields (for single node)**

Properties	Description
<b>General</b>	
<b>Name</b>	Enter the name of service graph template.
<b>Tenant</b>	To choose a tenant: <ol style="list-style-type: none"> <li>Click <b>Select Tenant</b>. The <b>Select Tenant</b> dialog appears.</li> <li>From the column on the left, click to choose a tenant.</li> <li>Click <b>Select</b>. You return to the <b>Create Service Graph</b> dialog box.</li> </ol>
<b>Description</b>	Enter a description of the service graph template.
<b>Settings</b>	
<b>Select a Device</b>	To choose a device: <ol style="list-style-type: none"> <li>Click <b>Select Device</b>. The <b>Select Device</b> dialog appears.</li> <li>From the column on the left, click to choose a device. Drag and drop the device in the <b>Drop Device</b> space below. This will open a small window where the actual device for this device type can be selected.</li> <li>Click <b>Select</b>. You return to the <b>Create Service Graph</b> dialog box.</li> </ol>

**Table 33: Create Service Graph Dialog Box Fields (for multinode)**

Properties	Description
<b>General</b>	
<b>Name</b>	Enter the name of service graph template.



Properties	Description
<b>Tenant</b>	To choose a tenant: <ol style="list-style-type: none"> <li>Click <b>Select Tenant</b>. The <b>Select Tenant</b> dialog appears.</li> <li>From the column on the left, click to choose a tenant.</li> <li>Click <b>Select</b>. You return to the <b>Create Service Graph</b> dialog box.</li> </ol>
<b>Description</b>	Enter a description of the service graph template.
<b>Settings:</b> Based on the required topology, drag and drop the devices in the box below	
<b>Application Load Balancer</b>	<ol style="list-style-type: none"> <li>Drag and drop the Application load balancer device into the box below.</li> <li>In the <b>Service node</b> dialog box, click on the <b>Select Application Load Balancer</b> and click to choose a Application Load Balancer in the left column then click <b>Add</b>.</li> </ol>
<b>Third Party Firewall</b>	<ol style="list-style-type: none"> <li>Drag and drop the Third Party Firewall next to the device in the box below.</li> <li>In the <b>Service node</b> dialog box, click on the <b>Third Party Firewall</b> and click to choose a Third Party Firewall in the left column then click <b>Add</b>.           <p><b>Note</b> Third Party Firewall cannot be the first node on the service graph.</p> </li> <li>If you want to enable the user-based redirect function on the <i>consumer</i> side of the Third Party Firewall, in the <b>Consumer Connector Type</b> field, place a check in the box next to the <b>Redirect</b> option.</li> <li>If you want to enable the user-based redirect function on the <i>provider</i> side of the Third Party Firewall, in the <b>Provider Connector Type</b> field, place a check in the box next to the <b>Redirect</b> option.</li> <li>In the <b>Provider Connector Type</b>, place a check next to the applicable option. Refer to <a href="#">About Layer 4 to Layer 7 Service Redirect</a> for information.</li> <li>Click <b>Add</b>.</li> </ol>
<b>Network Load Balancer</b>	<ol style="list-style-type: none"> <li>Drag and drop the Network load balancer device into the box below.</li> <li>In the <b>Service node</b> dialog box, click on the <b>Select Network Load Balancer</b> and click to choose a Network Load Balancer in the left column then click <b>Add</b>.</li> <li>If you want to enable the user-based redirect function on the <i>consumer</i> side of the network load balancer, in the <b>Consumer Connector Type</b> field, place a check in the box next to the <b>Redirect</b> option.</li> <li>If you want to enable the user-based redirect function on the <i>provider</i> side of the network load balancer, in the <b>Provider Connector Type</b> field, place a check in the box next to the <b>Redirect</b> option.</li> <li>Click <b>Add</b>.</li> </ol>

**Step 5** Click **Save** when finished.

**Step 6** The **EPG Communication** dialog box appears. Click on the **Go to details** to verify the Service Graph template.

---

## 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. This procedure is applicable for single node as well as multinode deployments.

### Before you begin

- You have configured the devices.
  - 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:

- a) Click **Select Contract**. The **Select Contract** dialog appears.
- b) 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:

- a) Click **Add Consumer EPGs**. The **Select Consumer EPGs** dialog appears.
- b) In the pane on the left side of the **Select Consumer EPGs** dialog, click the 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:

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

**Step 7** To choose a service graph:

- a) From the **EPG Communication Configuration** dialog, click **Select Service Graph**. The **Select Service Graph** dialog box appears.
- b) 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 34: Add Cloud Load Balancer Listener Dialog Box Fields For Application Gateway

Properties	Description
<b>Name</b>	Enter the name of the listener.
<b>Port</b>	Enter the port that the device will accept traffic on.
<b>Protocol</b>	For Application Gateway, click to choose <b>HTTP</b> or <b>HTTPS</b> .
<b>Security Policy</b>	Click the drop-down list and choose a security policy (only available when <b>HTTPS</b> is chosen).
<b>SSL Certificate</b>	<p>To choose an SSL certificate(only available when <b>HTTPS</b> is chosen):</p> <ol style="list-style-type: none"> <li>a. Click <b>Add SSL Certificates</b>.</li> <li>b. Click to place a check mark in the check box of the certificates you want to add.</li> <li>c. Choose a key ring: <ol style="list-style-type: none"> <li>1. Click <b>Select Key Ring</b>. The <b>Select Key Ring</b> dialog appears.</li> <li>2. From the <b>Select Key Ring</b> dialog, click to choose a key ring in the left column then click <b>Select</b>. The <b>Select Key Ring</b> dialog box closes.</li> </ol> </li> <li>d. Click the <b>Certificate Store</b> drop-down list and choose a certificate.</li> </ol> <p><b>Note</b> A listener can have multiple certificates.</p>
<b>Add Rule</b>	To add rule settings to the device listener, click <b>Add Rule</b> . A new row appears in the <b>Rules</b> list an the <b>Rules Settings</b> fields are enabled.

Properties	Description
<b>Rule Settings</b>	<p>The <b>Rule Settings</b> pane contains the following options:</p> <ul style="list-style-type: none"> <li>• <b>Name</b>—Enter a name for the rule.</li> <li>• <b>Host</b>—Enter a hostname to create a host-based condition. When a request is made for this hostname, the action you specify is taken.</li> <li>• <b>Path</b>—Enter a path to create a path-based condition. When a request is made for this path, the action you specify is taken.</li> <li>• <b>Type</b>—The action type tells the device which action to take. The action type options: <ul style="list-style-type: none"> <li>• <b>Return fixed response</b>—Returns a response using the following options: <ul style="list-style-type: none"> <li>• <b>Fixed Response Body</b>—Enter a response message.</li> <li>• <b>Fixed Response Code</b>—Enter a response code.</li> <li>• <b>Fixed response Content-Type</b>—Choose a content type.</li> </ul> </li> <li>• <b>Forward</b>—Forwards traffic using the following options: <ul style="list-style-type: none"> <li>• <b>Port</b>—Enter the port that the device will accept traffic on.</li> <li>• <b>Protocol</b>—Click to choose <b>HTTP</b> or <b>HTTPS</b>.</li> <li>• <b>Provider EPG</b>—The EPG with the web server that handles the traffic.</li> <li>• <b>EPG</b>—To choose an EPG: <ol style="list-style-type: none"> <li>a. Click <b>Select EPG</b>. The <b>Select EPG</b> dialog box appears.</li> <li>b. From the <b>Select EPG</b> dialog box, click to choose an EPG in the left column then click <b>Select</b>. The <b>Select EPG</b> dialog box closes.</li> </ol> </li> </ul> </li> <li>• <b>Redirect</b>—Redirects requests to another location using the following options: <ul style="list-style-type: none"> <li>• <b>Redirect Code</b>—Click the <b>Redirect Code</b> drop-down list and choose a code.</li> <li>• <b>Redirect Hostname</b>—Enter a hostname for the redirect.</li> <li>• <b>Redirect Path</b>—Enter a redirect path.</li> <li>• <b>Redirect Port</b>—Enter the port that the device will accept traffic on.</li> <li>• <b>Redirect Protocol</b>—Click to the <b>Redirect Protocol</b> drop-down list and choose <b>HTTP</b>, <b>HTTPS</b>, or <b>Inherit</b>.</li> <li>• <b>Redirect Query</b>—Enter a redirect query.</li> </ul> </li> </ul> </li> </ul>

Properties	Description
<b>Health Checks</b>	<p>The Application load balancer performs health checks on its backend pool targets for high availability. This can be configured under health checks:</p> <ul style="list-style-type: none"> <li>• <b>Protocol</b>-Click to choose <b>HTTP</b> or <b>HTTPS</b>.</li> <li>• <b>Path</b> - Enter the path. Default is /</li> <li>• <b>Port</b>-Enter a port on which health checks should be performed.</li> <li>• <b>Advanced Settings</b>- <ul style="list-style-type: none"> <li><b>Unhealthy Threshold</b>-Configure this threshold to determine when a backend target is advertised as unhealthy.</li> <li>• <b>Timeout</b> - Enter the value for health check timeout.</li> <li>• <b>Interval</b>-Enter a time in seconds to determine at what intervals checks should be performed.</li> <li>• <b>Success Code</b> - Enter the success code. Default is 200-399.</li> <li>• <b>Use host from rule</b> - Click on the checkbox if the hostname needs to be picked from the rule.</li> <li>• <b>Host</b> - If <b>Use host from rule</b> is not checked, provide the hostname to be used for health check.</li> </ul> </li> </ul> <p>Click <b>Add Rule</b> when finished.</p>

*Table 35: Add Cloud Load Balancer Listener Dialog Box Fields for Network Load Balancer*

Properties	Description
<b>Name</b>	Enter the name of the listener.
<b>Port</b>	Enter the port that the device will accept traffic on.
<b>Protocol</b>	Click to choose <b>TCP</b> or <b>UDP</b> .

Properties	Description
<b>Rule Settings</b>	<p>The <b>Rule Settings</b> pane contains the following options:</p> <ul style="list-style-type: none"> <li>• <b>Name</b>—Enter a name for the rule.</li> <li>• <b>Port</b>—Enter the port on which the backend pool servers will accept traffic from the load balancer.</li> <li>• <b>Protocol</b>-Click to choose <b>TCP</b> or <b>UDP</b>.</li> <li>• <b>Provider EPG</b>-The EPG with the web servers handling traffic.</li> <li>• <b>Type</b></li> <li>• <b>Forward</b>-The action type tells the device which action to take. The action type here is always <b>Forward</b>. Here the traffic is forwarded to the Port for EPG selected using the protocol chosen above.</li> <li>• <b>HA Port</b>- If you want to load balance traffic incoming on all the ports, instead of adding those many listeners a listener rule type ‘HA Ports’ can be configured for the same. This is a feature of <b>ONLY</b> the internal-facing load balancer.</li> </ul>
<b>Health Checks</b>	<p>The load balancer performs health checks on its backend pool targets for high availability. This can be configured here.</p> <ul style="list-style-type: none"> <li>• <b>Protocol</b>-Click to choose <b>TCP</b>, <b>HTTP</b> or <b>HTTPS</b>.</li> <li>• <b>Port</b>-Enter a port on which health checks should be performed.</li> <li>• <b>Advanced Settings</b>-</li> <li>• <b>Unhealthy Threshold</b>-Configure this threshold to determine when a backend target is advertised as unhealthy.</li> <li>• <b>Interval</b>-Enter a time in seconds to determine at what intervals checks should be performed.</li> </ul> <p>Click <b>Add Rule</b> when finished.</p>

**Step 10** Click **Add** when finished.  
The service graph is deployed.

## Deploying a Service Graph Using the REST API

The following sections describe how to deploy 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.

**Step 1** To create an internal-facing load balancer for Application Gateway:

**Example:**

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
  <fvTenant name="tn15">
    <fvRsCloudAccount tDn="uni/tn-infra/act-[<subscription id>-vendor-azure" />

    <cloudLB scheme="internal" type="application" name="alb-151-15" status="">
      <cloudRsLDevToCloudSubnet
tDn="uni/tn-tn15/ctxprofile-cProfilewestus15/cidr-[15.151.0.0/16]/subnet-[15.151.2.0/24]" />
      </cloudLB>

    </fvTenant>
  </polUni>
```

**Step 2** To create an internal-facing load balancer for Azure Load Balancing:

**Example:**

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
<fvTenant name="tn15">
<fvRsCloudAccount tDn="uni/tn-infra/act-[subscription id]-vendor-azure" />
<cloudLB scheme="internal" type="network" name="nlb-151-15" status="">
<cloudRsLDevToCloudSubnet
tDn="uni/tn-tn15/ctxprofile-cProfilewestus15/cidr-[15.151.0.0/16]/subnet-[15.151.2.0/24]" />
</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.

**Step 1** To create an internet-facing load balancer for Application Gateway:

**Example:**

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
  <fvTenant name="tn15">
    <fvRsCloudAccount tDn="uni/tn-infra/act-[<subscription id>-vendor-azure" />

    <cloudLB scheme="internet" type="application" name="alb-151-15" status="">
```

```

        <cloudRsLDevToCloudSubnet
tDn="uni/tn-tn15/ctxprofile-cProfilewestus151/cidr-[15.151.0.0/16]/subnet-[15.151.2.0/24]" />
        </cloudLB>

    </fvTenant>
</polUni>

```

**Step 2** To create an internet-facing load balancer for Azure Load Balancing:

**Example:**

```

<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
<fvTenant name="tn15">
<fvRsCloudAccount tDn="uni/tn-infra/act-[[subscription id]]-vendor-azure" />
<cloudLB scheme="internet" type="network" name="nlb-151-15" status="">
<cloudRsLDevToCloudSubnet
tDn="uni/tn-tn15/ctxprofile-cProfilewestus151/cidr-[15.151.0.0/16]/subnet-[15.151.2.0/24]" />
</cloudLB>
</fvTenant>
</polUni>

```

## Creating a Third-Party Firewall Using the REST API

This example demonstrates how to create a third-party firewall using the REST API.

This example demonstrates how to create a third-party firewall using the REST API:

**Example:**

```

<cloudLDev name="HubFW" svcType="FW" status="">
  <cloudRsLDevToCtx tDn="uni/tn-infra/ctx-overlay-2" />
  <cloudLIf name="provider">
    <cloudEPSelector name="east" matchExpression="IP=='[[eastus_FwUntrustSubnet]]'" status="" />
  </cloudLIf>
  <cloudLIf name="consumer">
    <cloudEPSelector name="east" matchExpression="IP=='[[eastus_FwTrustSubnet]]'" status="" />
  </cloudLIf>
</cloudLDev>

```



## Creating a Service Graph Using the REST API

This example demonstrates how to create a service graph using the REST API.

### Step 1 To create a service graph for Application Gateway:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
  <fvTenant name="tn15">

    <vnsAbsGraph name="c15_g1" type="cloud" status="">
      <vnsAbsTermNodeProv name="p1">
        <vnsAbsTermConn/>
      </vnsAbsTermNodeProv>
      <vnsAbsTermNodeCon name="c1">
        <vnsAbsTermConn/>
      </vnsAbsTermNodeCon>
      <vnsAbsNode managed="yes" name="N1" funcType="GoTo">
        <vnsRsNodeToCloudLDev tDn="uni/tn-tn15/clb-alb-151-15"/>
        <vnsAbsFuncConn name="provider"/>
        <vnsAbsFuncConn name="consumer"/>
      </vnsAbsNode>
      <vnsAbsConnection connDir="consumer" connType="external" name="con1">
        <vnsRsAbsConnectionConns tDn="uni/tn-tn15/AbsGraph-c15_g1/AbsTermNodeCon-c1/AbsTConn"/>
        <vnsRsAbsConnectionConns tDn="uni/tn-tn15/AbsGraph-c15_g1/AbsNode-N1/AbsFConn-consumer"/>
      </vnsAbsConnection>
      <vnsAbsConnection connDir="provider" connType="internal" name="con2">
        <vnsRsAbsConnectionConns tDn="uni/tn-tn15/AbsGraph-c15_g1/AbsTermNodeProv-p1/AbsTConn"/>
        <vnsRsAbsConnectionConns tDn="uni/tn-tn15/AbsGraph-c15_g1/AbsNode-N1/AbsFConn-provider"/>
      </vnsAbsConnection>
    </vnsAbsGraph>

  </fvTenant>
</polUni>
```

### Step 2 To create a service graph for Azure Load Balancing:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
  <fvTenant name="tn15">

    <vnsAbsGraph name="c15_g1" type="cloud" status="">

      <vnsAbsTermNodeProv name="p1">

        <vnsAbsTermConn />

      </vnsAbsTermNodeProv>

      <vnsAbsTermNodeCon name="c1">

        <vnsAbsTermConn />

      </vnsAbsTermNodeCon>

      <vnsAbsNode managed="yes" name="N1" funcType="GoTo">

        <vnsRsNodeToCloudLDev tDn="uni/tn-tn15/clb-nlb-151-15" />

      </vnsAbsNode>

    </vnsAbsGraph>

  </fvTenant>
</polUni>
```

```

<vnsAbsFuncConn name="provider" />
<vnsAbsFuncConn name="consumer" />
</vnsAbsNode>

<vnsAbsConnection connDir="consumer" connType="external" name="con1">
<vnsRsAbsConnectionConns tDn="uni/tn-tn15/AbsGraph-c15_g1/AbsTermNodeCon-c1/AbsTConn" />
<vnsRsAbsConnectionConns tDn="uni/tn-tn15/AbsGraph-c15_g1/AbsNode-N1/AbsFConn-consumer" />
</vnsAbsConnection>

<vnsAbsConnection connDir="provider" connType="internal" name="con2">
<vnsRsAbsConnectionConns tDn="uni/tn-tn15/AbsGraph-c15_g1/AbsTermNodeProv-p1/AbsTConn" />
<vnsRsAbsConnectionConns tDn="uni/tn-tn15/AbsGraph-c15_g1/AbsNode-N1/AbsFConn-provider" />
</vnsAbsConnection>

</vnsAbsGraph>
</fvTenant>
</polUni>

```

## Creating a Multi-Node Service Graph Using the REST API

This example demonstrates how to create a multi-node service graph using the REST API.

To create a multi-node service graph, enter a post such as the following example;

```

<polUni>
  <fvTenant name="tn12_iar_iavpc" status="">
    <fvRsCloudAccount tDn="uni/tn-infra/[SubscriptionId]-vendor-azure"/>
    <fvCtx name="vrf50" status="" />
    <fvCtx name="vrf60" status="" />
    <cloudVpnGwPol name="VgwPol0"/>
    <cloudCtxProfile name="c50" status="">
      <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-azure/region-westus"/>
      <cloudRsToCtx tnFvCtxName="vrf50"/>
      <cloudRsCtxProfileToGatewayRouterP tDn="uni/tn-infra/gwrouterp-default" status="" />
      <cloudCidr addr="12.3.0.0/16" primary="true" status="">
        <cloudSubnet ip="12.3.30.0/24" status="" name="GatewaySubnet" usage="gateway">
          <cloudRsZoneAttach tDn="uni/clouddomp/provp-azure/region-westus/zone-default"/>
        </cloudSubnet>
        <cloudSubnet ip="12.3.2.0/24" status="" name="ALBSubnet">
          <cloudRsZoneAttach tDn="uni/clouddomp/provp-azure/region-westus/zone-default"/>
        </cloudSubnet>
        <cloudSubnet ip="12.3.1.0/24" status="" name="FwMgmtSubnet">
          <cloudRsZoneAttach tDn="uni/clouddomp/provp-azure/region-westus/zone-default"/>
        </cloudSubnet>
        <cloudSubnet ip="12.3.3.0/24" status="" name="FwUntrustSubnet">
          <cloudRsZoneAttach tDn="uni/clouddomp/provp-azure/region-westus/zone-default"/>
        </cloudSubnet>
        <cloudSubnet ip="12.3.4.0/24" status="" name="FwTrustSubnet">

```

```

    <cloudRsZoneAttach tDn="uni/clouddomp/provp-azure/region-westus/zone-default"/>
  </cloudSubnet>
  <cloudSubnet ip="12.3.5.0/24" status="" name="ConsumerSubnet">
    <cloudRsZoneAttach tDn="uni/clouddomp/provp-azure/region-westus/zone-default"/>
  </cloudSubnet>
</cloudCidr>
</cloudCtxProfile>
<cloudCtxProfile name="c60" status="">
  <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-azure/region-westus2"/>
  <cloudRsToCtx tnFvCtxName="vrf60"/>
  <cloudRsCtxProfileToGatewayRouterP tDn="uni/tn-infra/gwrouterp-default" status="">
  <cloudCidr addr="12.4.0.0/16" primary="true" status="">
    <cloudSubnet ip="12.4.1.0/24" status="" name="ProviderSubnet">
      <cloudRsZoneAttach tDn="uni/clouddomp/provp-azure/region-westus2/zone-default"/>
    </cloudSubnet>
    <cloudSubnet ip="12.4.2.0/24" status="" name="NLBSubnet">
      <cloudRsZoneAttach tDn="uni/clouddomp/provp-azure/region-westus2/zone-default"/>
    </cloudSubnet>
    <cloudSubnet ip="12.4.30.0/24" status="" name="GatewaySubnet" usage="gateway">
      <cloudRsZoneAttach tDn="uni/clouddomp/provp-azure/region-westus2/zone-default"/>
    </cloudSubnet>
  </cloudCidr>
</cloudCtxProfile>
<cloudApp name="ap50" status="">
  <cloudEPg name="ap50vrf50epg1" status="">
    <cloudRsCloudEPgCtx tnFvCtxName="vrf50"/>
    <fvRsCons tnVzBrCPName="con50"/>
    <fvRsProv tnVzBrCPName="con60"/>
    <cloudEPSelector matchExpression="IP=='12.3.5.0/24'" name="100"/>
  </cloudEPg>
  <cloudEPg name="ap50vrf50epg2" status="">
    <cloudRsCloudEPgCtx tnFvCtxName="vrf50"/>
    <fvRsProv tnVzBrCPName="con60"/>
    <cloudEPSelector matchExpression="IP=='12.3.1.0/24'" name="100"/>
  </cloudEPg>
  <cloudExtEPg routeReachability="internet" name="ap50extepg1">
    <cloudExtEPSelector name="1" subnet="0.0.0.0/0"/>
    <cloudRsCloudEPgCtx tnFvCtxName="vrf50"/>
    <fvRsCons tnVzBrCPName="con60"/>
  </cloudExtEPg>
</cloudApp>
<cloudApp name="ap60" status="">
  <cloudEPg name="ap60vrf60epg1" status="">
    <cloudRsCloudEPgCtx tnFvCtxName="vrf60"/>
    <fvRsProv tnVzBrCPName="con50"/>
    <fvRsProv tnVzBrCPName="con70"/>
    <cloudEPSelector matchExpression="IP=='12.4.1.0/24'" name="100"/>
  </cloudEPg>
  <cloudExtEPg routeReachability="internet" name="ap60extepg1">
    <cloudExtEPSelector name="1" subnet="0.0.0.0/0"/>
    <cloudRsCloudEPgCtx tnFvCtxName="vrf60"/>
    <fvRsCons tnVzBrCPName="con70"/>
  </cloudExtEPg>
</cloudApp>
<vzBrCP name="con50" scope="tenant" status="">
  <vzSubj name="con50">
    <vzRsSubjFiltAtt tnVzFilterName="f10"/>
    <vzRsSubjGraphAtt tnVnsAbsGraphName="g1" status="">
  </vzSubj>
</vzBrCP>
<vzBrCP name="con60" scope="tenant" status="">
  <vzSubj name="con60">
    <vzRsSubjFiltAtt tnVzFilterName="f20"/>
  </vzSubj>

```

```

</vzBrCP>
<vzBrCP name="con70" scope="context" status="">
  <vzSubj name="con70">
    <vzRsSubjFiltAtt tnVzFilterName="f20"/>
  </vzSubj>
</vzBrCP>
<vzFilter name="f10" status="">
  <vzEntry etherT="ip" prot="icmp" name="f10entry1" status=""/>
  <vzEntry etherT="ip" prot="udp" dFromPort="1" dToPort="65535" name="f10entry2" status=""/>
  <vzEntry etherT="ip" prot="tcp" dFromPort="1" dToPort="65535" name="f10entry3" status=""/>
</vzFilter>
<vzFilter name="f20" status="">
  <vzEntry etherT="ip" prot="tcp" dFromPort="http" dToPort="http" name="f20entry1" status=""/>
  <vzEntry etherT="ip" prot="tcp" dFromPort="https" dToPort="https" name="f20entry2" status=""/>
  <vzEntry etherT="ip" prot="tcp" dFromPort="22" dToPort="22" name="f20entry3" status=""/>
</vzFilter>
<cloudLB name="FrontALB" type="application" scheme="internal" >
  <cloudRsLDevToCloudSubnet
tDn="uni/tn-tn12_iar_iavpc/ctxprofile-c50/cidr-[12.3.0.0/16]/subnet-[12.3.2.0/24]"/>
  </cloudLB>
<cloudLDev name="FW" svcType="FW" status="">
  <cloudRsLDevToCtx tDn="uni/tn-tn12_iar_iavpc/ctx-vrf50" />
  <cloudLif name="provider" >
    <cloudEPSelector name="1" matchExpression="custom:tagp=='trustFW'"/>
  </cloudLif>
  <cloudLif name="consumer" >
    <cloudEPSelector name="1" matchExpression="custom:tagp=='untrustFW'"/>
  </cloudLif>
</cloudLDev>
<cloudLB name="BackNLB" type="network" scheme="internal" >
  <cloudRsLDevToCloudSubnet
tDn="uni/tn-tn12_iar_iavpc/ctxprofile-c60/cidr-[12.4.0.0/16]/subnet-[12.4.2.0/24]"/>
  </cloudLB>
<vnsAbsGraph name="g1" type="cloud" status="" >
  <vnsAbsTermNodeProv name="Input1" >
    <vnsAbsTermConn name="C1"/>
  </vnsAbsTermNodeProv>
  <vnsAbsTermNodeCon descr="" name="Output1" nameAlias="" ownerKey="" ownerTag="">
    <vnsAbsTermConn name="C2" />
  </vnsAbsTermNodeCon>
  <vnsAbsNode funcType="GoTo" name="N1" managed="yes" funcTemplateType="ADC_ONE_ARM" >
    <vnsRsNodeToCloudLDev tDn="uni/tn-tn12_iar_iavpc/clb-FrontALB" />
    <vnsAbsFuncConn attNotify="no" descr="" name="provider" nameAlias="" ownerKey="" ownerTag=""/>
    <vnsAbsFuncConn attNotify="no" descr="" name="consumer" nameAlias="" ownerKey="" ownerTag=""/>
    <cloudSvcPolicy tenantName="tn12_iar_iavpc" contractName="con50" subjectName="con50" >
      <cloudListener name="http_listener1" port="80" protocol="http">
        <cloudListenerRule name="rule1" priority="20" default="yes" >
          <cloudRuleAction type="forward" port="80" protocol="http"/>
        </cloudListenerRule>
      </cloudListener>
    </cloudSvcPolicy>
  </vnsAbsNode>
  <vnsAbsNode funcType="GoTo" name="N2" managed="no" funcTemplateType="ADC_TWO_ARM" >
    <vnsRsNodeToCloudLDev tDn="uni/tn-tn12_iar_iavpc/cld-FW" />
    <vnsAbsFuncConn attNotify="no" descr="" connType="snat_dnat" name="provider" nameAlias=""
ownerKey="" ownerTag=""/>
    <vnsAbsFuncConn attNotify="no" descr="" connType="none" name="consumer" nameAlias="" ownerKey=""
ownerTag=""/>
  </vnsAbsNode>
  <vnsAbsNode funcType="GoTo" name="N3" managed="yes" funcTemplateType="ADC_ONE_ARM" >
    <vnsRsNodeToCloudLDev tDn="uni/tn-tn12_iar_iavpc/clb-BackNLB" />
    <vnsAbsFuncConn attNotify="no" descr="" name="provider" nameAlias="" ownerKey="" ownerTag=""/>
    <vnsAbsFuncConn attNotify="no" descr="" name="consumer" nameAlias="" ownerKey="" ownerTag=""/>
    <cloudSvcPolicy tenantName="tn12_iar_iavpc" contractName="con50" subjectName="con50" >

```

```

    <cloudListener name="http_listener1" port="80" protocol="tcp">
      <cloudListenerRule name="rule1" priority="20" default="yes" >
        <cloudRuleAction type="forward" port="80" protocol="tcp"
epgdn="uni/tn-tn12_iar_iavpc/cloudapp-ap60/cloudepg-ap60vrf60epg1"/>
        </cloudListenerRule>
      </cloudListener>
    </cloudSvcPolicy>
  </vnsAbsNode>
  <vnsAbsConnection connDir="provider" connType="external" name="CON4">
    <vnsRsAbsConnectionConns tDn="uni/tn-tn12_iar_iavpc/AbsGraph-g1/AbsNode-N3/AbsFConn-provider"/>
    <vnsRsAbsConnectionConns tDn="uni/tn-tn12_iar_iavpc/AbsGraph-g1/AbsTermNodeProv-Input1/AbsTConn"/>
  </vnsAbsConnection>
  <vnsAbsConnection connDir="consumer" connType="external" name="CON1">
    <vnsRsAbsConnectionConns tDn="uni/tn-tn12_iar_iavpc/AbsGraph-g1/AbsTermNodeCon-Output1/AbsTConn"/>
    <vnsRsAbsConnectionConns tDn="uni/tn-tn12_iar_iavpc/AbsGraph-g1/AbsNode-N1/AbsFConn-consumer"/>
  </vnsAbsConnection>
  <vnsAbsConnection connDir="consumer" connType="external" name="CON2">
    <vnsRsAbsConnectionConns tDn="uni/tn-tn12_iar_iavpc/AbsGraph-g1/AbsNode-N1/AbsFConn-provider"/>
    <vnsRsAbsConnectionConns tDn="uni/tn-tn12_iar_iavpc/AbsGraph-g1/AbsNode-N2/AbsFConn-consumer"/>
  </vnsAbsConnection>
  <vnsAbsConnection connDir="consumer" connType="external" name="CON3">
    <vnsRsAbsConnectionConns tDn="uni/tn-tn12_iar_iavpc/AbsGraph-g1/AbsNode-N2/AbsFConn-provider"/>
    <vnsRsAbsConnectionConns tDn="uni/tn-tn12_iar_iavpc/AbsGraph-g1/AbsNode-N3/AbsFConn-consumer"/>
  </vnsAbsConnection>
</vnsAbsGraph>
</fvTenant>
</polUni>

```

## Creating a Multi-Node Service Graph With Redirect Using the REST API

This example demonstrates how to create a multi-node service graph with redirect using the REST API.

### Step 1 To set up the infra tenant:

```

<polUni>
  <fabricInst>
    <commPol name="default">
      <commSsh name="ssh" adminSt="enabled" passwordAuth="enabled" />
    </commPol>
    <dnsProfile name="default">
      <dnsProv addr="172.23.136.143" preferred="yes" status=""/>
    </dnsProfile>
  </fabricInst>
  <fvTenant name="infra">
    <fvRsCloudAccount tDn="uni/tn-infra/act-[[{subscriptionId}]]-vendor-azure"/>
    <cloudAccount name="insbu" id="[[{subscriptionId}]]" vendor="azure" accessType="credentials"
status="">
      <cloudRsCredentials tDn="uni/tn-infra/credentials-cApicApp"/>
    </cloudAccount>
    <cloudCredentials name="cApicApp" keyId="[[{accessKeyId}]]" key="[[{accessKey}]]" httpProxy="">
      <cloudRsAD tDn="uni/tn-infra/ad-[[{adId}]]"/>
    </cloudCredentials>
    <cloudAD name="CiscoINSBUAd" id="[[{adId}]]" />
    <cloudApicSubnetPool subnet="10.10.1.0/24" />
    <cloudtemplateInfraNetwork name="default" numRoutersPerRegion="2" vrfName="overlay-1"
numRemoteSiteSubnetPool="1" status="">
      <cloudtemplateProfile name="default" routerUsername="cisco" routerPassword="ins3965" />
    </cloudtemplateInfraNetwork>
  </fvTenant>
</polUni>

```

```

    <cloudtemplateExtSubnetPool subnetpool="11.11.0.0/16" status=""/>
    <cloudtemplateExtNetwork name="default" status="">
      <cloudRegionName provider="azure" region="{{region}}"/>
      <cloudtemplateVpnNetwork name="default">
        <cloudtemplateIpSecTunnel peeraddr="{{peerAddress}}"/>
        <cloudtemplateOspf area="0.0.0.1" />
      </cloudtemplateVpnNetwork>
    </cloudtemplateExtNetwork>
    <cloudtemplateIntNetwork name="default">
      <cloudRegionName provider="azure" region="{{region}}"/>
    </cloudtemplateIntNetwork>
  </cloudtemplateInfraNetwork>
</fvTenant>
<cloudDomP>
  <cloudBgpAsP asn="1111"/>
  <cloudProvP vendor="azure">
    <cloudRegion adminSt="managed" name="{{region}}">
      <cloudZone name="default"/>
    </cloudRegion>
  </cloudProvP>
</cloudDomP>
</polUni>

```

## Step 2 To configure the service device in the hub VNet:

```

<polUni>
  <fvTenant name="infra">
    <fvRsCloudAccount tDn="uni/tn-infra/act-[[subscriptionId]]-vendor-azure"/>
    <cloudCtxProfile name="ct_ctxprofile_{{region}}" status="modified">
      <cloudRsCtxProfileToRegion status="" tDn="uni/clouddomp/provp-azure/region-{{region}}"/>

      <cloudCidr name="cidr1" addr="{{HubCidrSvc}}" primary="no" status="">
        <cloudSubnet ip="{{HubNLBSubnet}}" name="HubNLBSubnet" status="">
          <cloudRsZoneAttach status=""
tDn="uni/clouddomp/provp-azure/region-{{region}}/zone-default"/>
        </cloudSubnet>
        <cloudSubnet ip="{{HubFWSubnetInt}}" name="HubFWSubnetInt" status="">
          <cloudRsZoneAttach status=""
tDn="uni/clouddomp/provp-azure/region-{{region}}/zone-default"/>
        </cloudSubnet>
        <cloudSubnet ip="{{HubFWSubnetExt}}" name="HubFWSubnetExt" status="">
          <cloudRsZoneAttach status=""
tDn="uni/clouddomp/provp-azure/region-{{region}}/zone-default"/>
        </cloudSubnet>
        <cloudSubnet ip="{{HubFWMgmtSubnet}}" name="HubFWMgmtSubnet" status="">
          <cloudRsZoneAttach status=""
tDn="uni/clouddomp/provp-azure/region-{{region}}/zone-default"/>
        </cloudSubnet>
        <cloudSubnet ip="{{ConsHubEPgSubnet}}" name="ConsHubEPgSubnet" status="">
          <cloudRsZoneAttach status=""
tDn="uni/clouddomp/provp-azure/region-{{region}}/zone-default"/>
        </cloudSubnet>
      </cloudCidr>
    </cloudCtxProfile>
    <cloudLDev name="{{FWName}}" status="">
      <cloudRsLDevToCtx tDn="uni/tn-infra/ctx-{{ServiceVNetName}}"/>
      <cloudLIf name="external" >
        <cloudEPSelector matchExpression="custom:EPG=='FwExt'" name="1"/>
      </cloudLIf>
      <cloudLIf name="internal" >
        <cloudEPSelector matchExpression="custom:EPG=='FwInt'" name="1"/>
      </cloudLIf>
    </cloudLDev>
  </fvTenant>
</polUni>

```

```

        <cloudLB name="{{NLBName}}" type="network" scheme="internal" size="small" instanceCount="2"
status="">
        <cloudRsLDevToCloudSubnet
tDn="uni/tn-infra/ctxprofile-ct_ctxprofile_{{region}}/cidr-{{HubCidrSvc}}/subnet-{{HubNLBSubnet}}">
        status=""/>
        </cloudLB>
    </fvTenant>
</polUni>

```

### Step 3 To configure a provider and the graph in a spoke:

```

<polUni>
    <fvTenant name="{{tnNameProv}}" status="" >
        <fvRsCloudAccount tDn="uni/tn-infra/act-{{subscriptionId}}-vendor-azure"/>
        <fvCtx name="{{ProviderVNetName}}"/>
        <cloudCtxProfile name="{{ProviderVNetName}}" status="">
            <cloudRsCtxProfileToGatewayRouterP tDn="uni/tn-infra/gwrouterp-default" status=""/>
            <cloudRsCtxProfileToRegion status="" tDn="uni/clouddomp/provp-azure/region-{{region}}"/>

            <cloudRsToCtx tnFvCtxName="{{ProviderVNetName}}"/>
            <cloudCidr name="cidr1" addr="{{VnetCidrProv}}" primary="yes" status="">
                <cloudSubnet ip="{{ProviderSubnet}}" name="ProviderSubnet" status="">
                    <cloudRsZoneAttach status=""
tDn="uni/clouddomp/provp-azure/region-{{region}}/zone-default"/>
                </cloudSubnet>
                <cloudSubnet ip="{{BackALBSubnet}}" name="BackALBSubnet" status="">
                    <cloudRsZoneAttach status=""
tDn="uni/clouddomp/provp-azure/region-{{region}}/zone-default"/>
                </cloudSubnet>
            </cloudCidr>
        </cloudCtxProfile>
        <!-- contract-->
        <vzFilter descr="" name="HttpsFilter" ownerKey="" ownerTag="">
            <vzEntry dFromPort="443" dToPort="443" etherT="ip" name="https" prot="tcp" status=""/>
            <vzEntry dFromPort="80" dToPort="80" etherT="ip" name="http" prot="tcp" status=""/>
            <vzEntry dFromPort="22" dToPort="22" etherT="ip" name="ssh" prot="tcp" status=""/>
        </vzFilter>
        <vzBrCP name="{{contractName}}" scope="global" status="">
            <vzSubj name="Sub1" revFltPorts="yes">
                <vzRsSubjGraphAtt directives="" tnVnsAbsGraphName="{{graphName}}"/>
                <vzRsSubjFiltAtt tnVzFilterName="HttpsFilter"/>
            </vzSubj>
        </vzBrCP>
        <!-- cloud App Profile-->
        <cloudApp name="provApp" status="">
            <cloudEPg name="App" status="">
                <cloudRsCloudEPgCtx tnFvCtxName="{{ProviderVNetName}}"/>
                <cloudEPSelector matchExpression="custom:EPG=='App'" name="1"/>
                <fvRsProv status="" tnVzBrCPName="{{contractName}}"/>
                <fvRsProv tnVzBrCPName="mgmt_common"/>
            </cloudEPg>
        </cloudApp>
        <!-- Abs Graph Creation -->
        <vnsAbsGraph name="{{graphName}}" uiTemplateType="UNSPECIFIED" type="cloud">
            <vnsAbsTermNodeProv name="T2">
                <vnsOutTerm/>
                <vnsInTerm />
                <vnsAbsTermConn attNotify="no" name="1" />
            </vnsAbsTermNodeProv>
            <vnsAbsTermNodeCon name="T1" >
                <vnsOutTerm/>
                <vnsInTerm />
                <vnsAbsTermConn attNotify="no" name="1" />
            </vnsAbsTermNodeCon>
        </vnsAbsGraph>
    </fvTenant>
</polUni>

```

```

</vnsAbsTermNodeCon>
<vnsAbsNode name="{NLBName}" managed="yes" >
  <vnsRsNodeToCloudLDev tDn="uni/tn-infra/clb-{{NLBName}}" status=""/>
  <cloudSvcPolicy tenantName="{{tnNameProv}}" contractName="{{contractName}}"
subjectName="Sub1" status="">
    <cloudHealthProbe name="http_listener1-rule1" protocol="tcp" port=22 interval=15
unhealthyThreshold=2/>
    <cloudListener name="http_listener1" port="80" protocol="tcp" status="">
      <cloudListenerRule name="rule1" default="true">
        <cloudRuleAction type="haPort" port="80" protocol="tcp"
healthProbe="http_listener1-rule1"/>
      </cloudListenerRule>
    </cloudListener>
  </cloudSvcPolicy>
  <vnsAbsFuncConn attNotify="no" name="provider" connType="redir"/>
  <vnsAbsFuncConn attNotify="no" name="consumer" connType="redir"/>
</vnsAbsNode>
<vnsAbsNode funcTemplateType="FW_ROUTED" name="{{FWName}}" managed="no">
  <vnsRsNodeToCloudLDev tDn="uni/tn-infra/cld-{{FWName}}" />
  <vnsAbsFuncConn attNotify="no" name="consumer" deviceLIIfName="internal"/>
  <vnsAbsFuncConn attNotify="no" name="provider" deviceLIIfName="internal"/>
</vnsAbsNode>
<vnsAbsNode name="{{BackALBName}}" managed="yes">
  <vnsRsNodeToCloudLDev tDn="uni/tn-{{tnNameProv}}/clb-{{BackALBName}}"/>
  <cloudSvcPolicy tenantName="{{tnNameProv}}" contractName="{{contractName}}"
subjectName="Sub1" status="">
    <cloudListener name="http_listener1" port="80" protocol="http" status="">
      <cloudListenerRule name="rule1" default="true">
        <cloudRuleAction type="forward" port="80" protocol="http"
epgdn="uni/tn-{{tnNameProv}}/cloudapp-provApp/cloudepg-App"/>
      </cloudListenerRule>
    </cloudListener>
  </cloudSvcPolicy>
  <vnsAbsFuncConn attNotify="no" name="provider"/>
  <vnsAbsFuncConn attNotify="no" name="consumer"/>
</vnsAbsNode>
<vnsAbsConnection adjType="L3" connDir="provider" connType="external" directConnect="no"
name="ConstermToNLB">
  <vnsRsAbsConnectionConns
tDn="uni/tn-{{tnNameProv}}/AbsGraph-{{graphName}}/AbsTermNodeCon-T1/AbsTConn"/>
  <vnsRsAbsConnectionConns
tDn="uni/tn-{{tnNameProv}}/AbsGraph-{{graphName}}/AbsNode-{{NLBName}}/AbsFConn-consumer"/>
</vnsAbsConnection>
<vnsAbsConnection adjType="L3" connDir="provider" connType="external" directConnect="no"
name="NLBtoFW">
  <vnsRsAbsConnectionConns
tDn="uni/tn-{{tnNameProv}}/AbsGraph-{{graphName}}/AbsNode-{{NLBName}}/AbsFConn-provider" />
  <vnsRsAbsConnectionConns
tDn="uni/tn-{{tnNameProv}}/AbsGraph-{{graphName}}/AbsNode-{{FWName}}/AbsFConn-consumer"/>
</vnsAbsConnection>
<vnsAbsConnection adjType="L3" connDir="provider" connType="external" directConnect="no"
name="FWtoBackALB">
  <vnsRsAbsConnectionConns
tDn="uni/tn-{{tnNameProv}}/AbsGraph-{{graphName}}/AbsNode-{{FWName}}/AbsFConn-provider" />
  <vnsRsAbsConnectionConns
tDn="uni/tn-{{tnNameProv}}/AbsGraph-{{graphName}}/AbsNode-{{BackALBName}}/AbsFConn-consumer"/>
</vnsAbsConnection>
<vnsAbsConnection adjType="L3" connDir="provider" connType="external" directConnect="no"
name="BackALBtoProv">
  <vnsRsAbsConnectionConns
tDn="uni/tn-{{tnNameProv}}/AbsGraph-{{graphName}}/AbsNode-{{BackALBName}}/AbsFConn-provider" />
  <vnsRsAbsConnectionConns
tDn="uni/tn-{{tnNameProv}}/AbsGraph-{{graphName}}/AbsTermNodeProv-T2/AbsTConn"/>
</vnsAbsConnection>

```



```

        </vnsAbsGraph>
        <cloudLB name="{{BackALBName}}" type="application" scheme="internal" size="small"
instanceCount="2">
            <cloudRsLDevToCloudSubnet
tDn="uni/tn-{{tnNameProv}}/ctxprofile-{{ProviderVNetName}}/cidr-{{VnetCidrProv}}/subnet-{{BackALBSubnet}}">
                status=""/>
            </cloudLB>
        </fvTenant>
</polUni>

```

**Step 4** To configure the consumer and import the contract defined in the provider:

```

<polUni>
    <fvTenant name="{{tnNameCons}}" >
        <fvRsCloudAccount tDn="uni/tn-infra/act-{{subscriptionId}}-vendor-azure"/>
        <fvCtx name="{{ConsumerVNetName}}"/>
        <cloudCtxProfile name="{{ConsumerVNetName}}" status="">
            <cloudRsCtxProfileToGatewayRouterP tDn="uni/tn-infra/gwrouterp-default" status=""/>
            <cloudRsCtxProfileToRegion status="" tDn="uni/cloudcomp/provp-azure/region-{{region}}"/>

            <cloudRsToCtx tnFvCtxName="{{ConsumerVNetName}}"/>
            <cloudCidr name="cidr1" addr="{{VnetCidrCons}}" primary="yes" status="">
                <cloudSubnet ip="{{ConsumerSubnet}}" name="ConsumerSubnet" status="">
                    <cloudRsZoneAttach status=""
tDn="uni/cloudcomp/provp-azure/region-{{region}}/zone-default"/>
                </cloudSubnet>
            </cloudCidr>
        </cloudCtxProfile>
        <vzCPIf name="imported_{{contractName}}">
            <vzRsIf tDn="uni/tn-{{tnNameProv}}/brc-{{contractName}}"/>
        </vzCPIf>
        <!-- cloud App Profile-->
        <cloudApp name="consApp" status="">
            <cloudEPg name="Web" status="">
                <cloudRsCloudEPgCtx tnFvCtxName="{{ConsumerVNetName}}"/>
                <cloudEPSelector matchExpression="custom:EPG=='Web'" name="1"/>
                <fvRsConsIf tnVzCPIfName="imported_{{contractName}}"/>
                <fvRsProv tnVzBrCPName="mgmt_common"/>
            </cloudEPg>
        </cloudApp>
    </fvTenant>
</polUni>

```

## Attaching a Service Graph Using the REST API

This example demonstrates how to attach a service graph using the REST API.

**Step 1** To attach a service graph for Application Gateways:

```

<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
    <fvTenant name="tn15">

        <vzBrCP name="c1">
            <vzSubj name="c1">
                <vzRsSubjGraphAtt tnVnsAbsGraphName="c15_g1"/>
            </vzSubj>
        </vzBrCP>
    </fvTenant>
</polUni>

```

```

    </vzBrCP>

  </fvTenant>
</polUni>

```

**Step 2** To attach a service graph for Azure Load Balancing:

```

<?xml version="1.0" encoding="UTF-8"?>

<!-- api/node/mo/uni/.xml -->

<polUni>

  <fvTenant name="tn15">

    <vzBrCP name="c1">

      <vzSubj name="c1">

        <vzRsSubjGraphAtt tnVnsAbsGraphName="c15_g1" />

      </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.

**Step 1** To create an HTTP service policy for Application Gateways:

```

<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>
          </cloudSvcPolicy>
        </vnsAbsNode>
      </vnsAbsGraph>
    </fvTenant>
  </polUni>

```

```

        </cloudListener>
    </cloudSvcPolicy>
</vnsAbsNode>
</vnsAbsGraph>
</fvTenant>
</polUni>

```

**Step 2** To create an HTTP service policy for Azure Load Balancing:

```

<?xml version="1.0" encoding="UTF-8"?>
<polUni>
<fvTenant name="tn15">
<vnsAbsGraph name="CloudGraph" type="cloud" status="">
<vnsAbsNode funcType="GoTo" name="N1" managed="yes">
<cloudSvcPolicy tenantName=" tn15" contractName="httpFamily" subjectName="consubj">
<cloudListener name="tcp_listener" port="80" protocol="tcp" status="">
<cloudListenerRule name="rule1" priority="10" default="yes" status="">
<cloudRuleAction type="forward" port="80" protocol="tcp" epgdn="uni/tn-
tn15/cloudapp-ap/cloudepg-provEPG" />
</cloudListenerRule>
</cloudListener>
<cloudListener name="udp_listener" port="55" protocol="udp" status="">
<cloudListenerRule name="rule1" priority="10" default="yes" status="">
<cloudRuleAction type="forward" port="55" protocol="udp" epgdn="uni/tn-
tn15/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*.



**Note** This procedure is applicable only for Application Gateways.

To configure a key ring:

```
<polUni>
  <fvTenant name="tn15" >
    <cloudCertStore>
      <pkiKeyRing status="" name="lbCert" tp="lbTP" key="-----BEGIN RSA PRIVATE KEY-----
MIIEpQIBAAKCAQEAA4DGxak+RHv/nToHLnmDBq2BfLimgX/zNJQC9bGuzr8Mj7dm0
XuHfQYGV0h1PtL4Pdx:5qjB0NbHjAVB1Gw8cDiErEgAXy9Km27ySo2foKryNQCRe
Ginn/CgF75QPied568eScNdZPt/eMeHAuRX/PykKUatWWncGanjvHqc+SOLPF6TD
gQ5nwOHfVvyM2DY8bfdYwRwMsO7JqZzbPMptA2QWblILsSoIrdkIlgf6ZfYy/EN
bH+nYN2rJT81zYsxz0YmR0oRQHTiN2NiDY/ZV63yxCXfLg9qpNZCuD8KOfdCZPEq
8takiWBxiR5/HRPscwAdWQsoiKgG1k4NEbFA9QIDAQABaoIBAQQDQQA9IS1YrdtqN
q6mZ3s2BNff/4kgb7gn0DWS+9EJJLcJNZVhFEo2ZxxxfPp6HRnjYS50W83/ElanD
+GD1bSutXuxqFWIQVh7r1ebYZIWk+NYSjr5yNVxux8U2hcNNV8WWVqkJjKcUqICB
Bm47FKj53LV46zeOgyCaibFrYxzJ9+farGneyBdnoV+3thmez7534Kci0t3J3Eri
lgSY3q16hPXB2ZXAP4jdAoLgWUDU4I1M6OqOiWopZM/QYIE/WtPYyJ0QzNCXObtC5
FboDcvedsgd4x5GLfV2A4xTBQMCTZUZJ9fYAcFogTZXD+UVqxorh47tf/mz+1fjq
flXphED1AoGBAPVlvKfGW46qqRnYovfryxxxz4OmlsVSGcJpQTQtBQI2koJ80weZJ
2s+CX0r+oDqwP23go/QEVYVkcic9RGkJBNGel+dm/bTjzgmQYtqSCNtecTsZD5JN
yljkciiiznDkjCjReSZ2kh3dGXlbrIYk7ezp2z7EKfDrHe5x5ouGMgCnaoGBAOnh
buDEohv8KJaB+DiUfhtoa3aKNPBO+zWPChp0HFGjPXshJcIYZc1GcycmuDKVnNd
MxhE/yOnQHowi4T9FMLpz5yh5zuCUVqOBgB1P6MzbC5t5MtLrEYr/AqFN1lCqyXQ
cVcT6icW1OAFJRw3c/OiESwLMzchsl8RnbwOi6kDaOGBANVlzmPb07zB3eGTCU0t
KGIqwfLncUkVaDZRZFZYPpNwiRkoe73j9brkNbgCqxW+NLp5UjoeFry0N6y106q/
ZA4I7FnXryLBw2HYuw41Vix1+XOZ/HeO3RmFN1z717dGmaGbv43aKIB9x+X5n8wF
6z1NtBhmBk7yNwomlIRagl5bAogAX0p4cJ/tJNXSe7AswHDQCL68uimJdDfZ5nKG
k83nE+Qc0qQozDJAmCisFmuSNRnSep3fiafjBFXK0X4h+mdbJCC7bagRnI92Mh0X
mOWsp4P2GdywkZwdbuHQ6UBp1Ferf9aztzTn+as6xKOUATEezy9DK9zMWzQhhtaY
m9yZTp0CgYEAlUtcpWjAzQbXODJGmxGdAAakPpeiKw/Da3MccrTdgJt88ezM10ej
Pdoab0G2PcfGjZotSGk7N4XARVkeq7pgZ0kwcYash06A2Hal+D1z/bGoZP+kmD/x
Ny82phxYXcnEc5Vv92lU59+j7e067UFLAYJeifuoFImvoFRnP4DIQ=
-----END RSA PRIVATE KEY-----" cert="-----BEGIN CERTIFICATE-----
MIIElTCCA32gAwIBAgIJAKWNjp//arBsMA0GCSqGSIb3DQEBCwUAMIGNQswCQYD
VQQTGEwJVVzELMAkGA1UECBMCMQ0ExETAPBqNVBacTCFNBhbiBkb3N1MRRlWfEAYDVQK
EwlNeUNvbXBhbnkxkjAMBGNVBAsTBU15T3JnMRgwFgYDVQQDFAs8qLmFtYXpvbmF3
cy5jb20xIDAeBgkqhkiG9w0BCQEWEWEXJhbXNoYWhAY2lZY28uY29tMB4XDTE4MTAw
MjIwMjIzY28uY29tggkApY20n/9qsGwwDAYDVROTBAUwAwEB/zANBgkqhkiG9w0B
AQsFAAOCAQEAE/RuzCheLiBhrurGet6eaVx9DPYydnIKVBSAKO+5iur84mQzhoT
nx5CN109xu5m15baCYZsSnn6D7usC092bPA/kRCGxt29gkjPWA74tJHQyHvwgBM
mOrLISHoelew+wRl0oVrChlTfKtXO68Tuk6vrqpW76hkFOIa7b2h1IIMdq6VA/
+A5FQ0xqYfKdVd2RaINpzI8mqZiszw+7E6j1PL5k44ftWEaYpGF1VesFEyJEL
gHBUiPt8TIbaMYI8qUQmB/emnLXekQ5PRxdRn1eA3h8jfQ3D1CQRTLmDL3tpFwg
qopM6et5ZKqShX4T87BsgZIOiquzXqsuHg==
```

```

-----END CERTIFICATE-----">
  </pkiKeyRing>

  <pkiTP status="" name="lbTP" certChain="-----BEGIN CERTIFICATE-----
MIIElTCCA32gAwIBAgIJAKWNjp//arBsMA0GCSqGSIB3DQEBCwUAMIGNMQswCQYD
VQQGEwJVUzELMAkGA1UECBMCQ0EwETAPBgNVBACTFCNhbiBkb3NlMRlWEAYDVQK
EwlnEUNvbXBhbnkxDjAMBGNVBAStBU15T3JnMRgwFgYDVQQDFA8qLmFtYXpvc2F3
cy5jb20xIDAeBgkqhkiG9w0BCQEWEXJhbXNoYWhAY21zY28uY29tMB4XDTE4MTAw
MjIwNTMwNV0xMjE5MjAwMjIwNTMwNVowY0x0c2ZAJBgNVBAYTA1VTMzQYDVQKI
EwJQTERMA8GA1UEBxMIU2FuIEpvc2UxZjAQBGNVBAoTCU15Q292tCGFueTEOMAww
A1UECzMTXlPcmcxGDAWBgNVBAMUDyouYWh1hem9uYXZzLmNvbTEgMB4GCSqGSIB3
DQEJARYRcmFtc2hhaEBjaXNjby5jb20wggeiMa0GCSqGSIB3DQEBAQUAA4IBDwAw
ggEKAoIBAQQDgMbFor5Ee/+dOgcueYMGryF8uKaBf/M01AL1sa70vwyPt2bRe4d9B
ga/SHU+0vg93F/mqMHQ1seMBUHUbDxwOISsSABfL0qbbvJKjZ+gqvI2oJF4aKef8
KAXv1A8h53nrx5Jw0Nk+394x4cC5Ff8/KQpRq1ZadwZqe08epz5I4s8XpMOBDMfa
4ccW/IzYNjxt9lhataYaw7smpnNs8ym0DZBZuUguxKgit2QgiB/p19jL8Q1sf6dg
3aslPyXNizHPRIzHSHfAdOI3Y2INj9lXrfLEJd8uD2qk1kK4Pwo590Jk8Sry1qSJ
YHGJHn8dE+xxYB1ZCyIqAbWTg0RsUD1AgMBAAGjgfUwgfIwhQYDVR0OBByEFBYq
K3b39+1oOr4IBSsePwcOpML7MIHCBgNVHSMEgbowgbeAFBYqK3b39+1oOr4IBSse
PwcOpML7oYGTpIGQMIGNMQswCQYDVQGEwJVUzELMAkGA1UECBMCQ0EwETAPBgNV
BACTFCNhbiBkb3NlMRlWEAYDVQKKEwlnEUNvbXBhbnkxDjAMBGNVBAStBU15T3Jn
MRgwFgYDVQQDFA8qLmFtYXpvc2F3cy5jb20xIDAeBgkqhkiG9w0BCQEWEXJhbXNo
YWhAY21zY28uY29tggkApY2On/9qsGwwDAYDVR0TBAUwAwEB/zANBgkqhkiG9w0B
AQsFAAOCAQEAE/RuzCheLIbHbrurGet6eaVx9DPYydNiKVBSAKO+5iuR84mQzhoT
nx5CN109xu5ml5baCYZzSnn6D7usC092bPA/kRCGxt29gkjpWA74tJHqIhVWgbM
mOrLiShoeIewv+wR10oVRChlTfKtXO68Tuk6vrqpW76hKfOHia7b2h1IIMdq6VA/
+A5FQ0xqYfKdVd2RaINpzI8mqZiszqw+7E6j1PL5k4tftWEaYpfGP1VesFEyJEL
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.



**Note** This is applicable only for the Application Gateways.

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="default"
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>
            </cloudListener>
          </cloudSvcPolicy>
        </vnsAbsNode>
      </vnsAbsGraph>
    </fvTenant>
  </polUni>
```



## CHAPTER

# 7

# Cisco Cloud APIC Security

---

This chapter contains the following sections:

- [Access, Authentication, and Accounting, on page 147](#)
- [Configuring TACACS+, RADIUS, LDAP and SAML Access, on page 148](#)
- [Configuring HTTPS Access, on page 155](#)

## 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 69](#) to configure a Local User and associate it to the OTP, SSH Public Key, and X.509 User Certificate using the 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 Cloud APIC.

## Overview

This topic provides step-by-step instructions on how to enable access to the 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**.

- On the menu bar, choose **Administrative > Authentication**.
- In the Work pane, click on **Providers** tab and then click on the **Actions** drop-down and select **Create Provider**.  
The **Create Provider** dialog box appears.
- In the **Host Name/IP Address** field, enter the Host Name/IP Address of the provider.
- In the **Description** field, enter a description of the provider.
- Click the **Type** drop-down list and choose **TACACS+**.
- 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+.

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



Properties	Description
<b>General</b>	
<b>Name</b>	Enter the name of the Login Domain
Description	Enter the description of the Login Domain.
<b>Settings</b>	
Realm	Choose <b>TACACS+</b> from the dropdown menu
Providers	To choose a Provider(s): <ol style="list-style-type: none"> <li>1. Click <b>Add Providers</b>. The <b>Select Providers</b> dialog appears.</li> <li>2. Click to choose a provider(s) in the column on the left.</li> <li>3. Click <b>Select</b>. You return to the <b>Create Login Domain</b> dialog box.</li> </ol>

- 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**.

- On the menu bar, choose **Administrative > Authentication**.
- In the Work pane, click on **Providers** tab and then click on the **Actions** drop-down and select **Create Provider**.

The **Create Provider** dialog box appears.

- In the **Host Name/IP Address** field, enter the Host Name/IP Address of the provider.
- In the **Description** field, enter a description of the provider.
- 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 **Intent** icon.

The **Intent** menu appears.

- b) Click the drop-down arrow below the **Intent** search box and choose **Administrative**

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

- c) From the **Administrative** list in the **Intent** 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
<b>General</b>	
<b>Name</b>	Enter the name of the Login Domain
Description	Enter the description of the Login Domain.
<b>Settings</b>	
Realm	Choose <b>RADIUS</b> from the dropdown menu
Providers	<p>To choose a Provider(s):</p> <ol style="list-style-type: none"> <li>1. Click <b>Add Providers</b>. The <b>Select Providers</b> dialog appears.</li> <li>2. Click to choose a provider(s) in the column on the left.</li> <li>3. Click <b>Select</b>. You return to the <b>Create Login Domain</b> dialog box.</li> </ol>

- 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**.

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 **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
<b>General</b>	
<b>Name</b>	Enter the name of the Login Domain
Description	Enter the description of the Login Domain.
<b>Settings</b>	
Realm	Choose <b>LDAP</b> from the dropdown menu
Providers	<p>To choose a Provider(s):</p> <ol style="list-style-type: none"> <li>1. Click <b>Add Providers</b>. The <b>Select Providers</b> dialog appears.</li> <li>2. Click to choose a provider(s) in the column on the left.</li> <li>3. Click <b>Select</b>. You return to the <b>Create Login Domain</b> dialog box.</li> </ol>

Properties	Description
Authentication Type	<ol style="list-style-type: none"> <li>1. Select <b>Cisco AV Pairs</b>, if provider(s) was configured with <b>CiscoAVPair</b> as the <b>Attribute</b>.</li> <li>2. Select <b>LDAP Group Map Rules</b>, if provider(s) was configured with <b>memberOf</b> as the <b>Attribute</b>. <ol style="list-style-type: none"> <li>a. Click <b>Add LDAP Group Map Rule</b>. The dialog box appears.</li> <li>b. Specify the map rule <b>Name</b>, <b>Description</b> (optional), and <b>Group DN</b>.</li> <li>c. Click the + next to <b>Add Security Domain</b>. The dialog box appears.</li> <li>d. Click the + to access the <b>Role</b> name and <b>Role Privilege</b> Type (<b>Read</b> or <b>Write</b>) fields. Click check mark.</li> <li>e. Repeat step 4 to add more roles. Then click <b>Add</b>.</li> <li>f. Repeat step 3 to add more security domains. Then click <b>Add</b>.</li> </ol> </li> </ol>

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
<b>General</b>	
<b>Name</b>	Enter the name of the Login Domain
Description	Enter the description of the Login Domain.
<b>Settings</b>	
Realm	Choose <b>SAML</b> from the dropdown menu
Providers	To choose a Provider(s): <ol style="list-style-type: none"> <li>1. Click <b>Add Providers</b>. The <b>Select Providers</b> dialog appears.</li> <li>2. Click to choose a provider(s) in the column on the left.</li> <li>3. Click <b>Select</b>. You return to the <b>Create Login Domain</b> dialog box.</li> </ol>

- 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, enter a name for the key ring in the **Name** field and a description in the **Description** field.
- 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 8

# Configuration Drifts

---

- [Configuration Drift Notifications and Faults, on page 159](#)
- [Enabling Configuration Drift Detection, on page 160](#)
- [Checking for Missing Contracts Configuration, on page 161](#)
- [Configuration Drift Troubleshooting, on page 164](#)

## 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. Future releases will provide the configuration drift visibility into the other Cloud APIC objects as well as information about extraneous configurations deployed in the cloud but not defined in the Cloud APIC.

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?

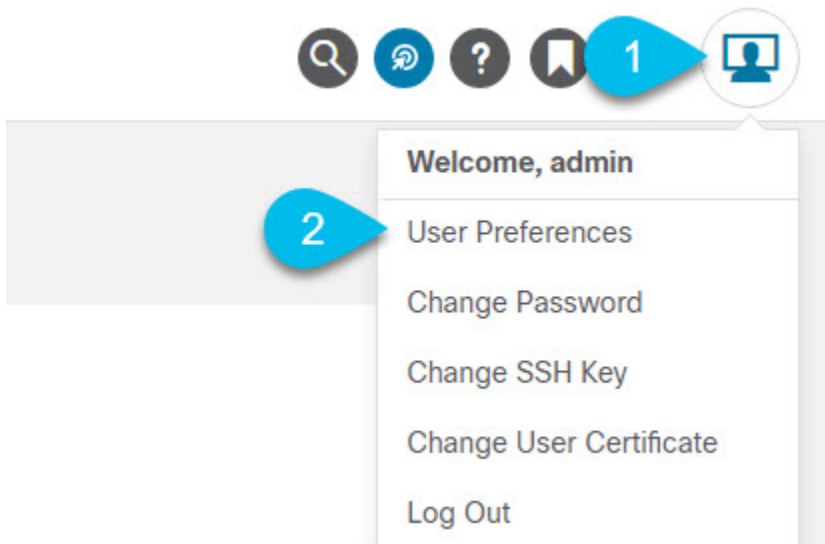
Similarly to the previous scenario, this can occur if there are connection or API issues or if a cloud administrator manually creates additional security rules directly in the cloud provider's UI. Any existing but not intended configuration may present issues.

# Enabling Configuration Drift Detection

In this release, configuration drift detection is in beta stage, as such it is disabled by default. This section describes how to enable configuration drift detection in your Cloud APIC user preferences.


**Step 1** Log in to your Cloud APIC GUI.

**Step 2** Open the **User Preferences** dialog.



- a) In the top right corner of the screen, click the user icon.
- b) From the menu, select **User Preferences**.

**Step 3** In the **User Preferences** dialog, enable **Configuration Drift Detection**.



- a) Check the **Enabled** checkbox.
- b) Click **Done** to save the change.

---

## 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** Navigate to the **Configuration Drifts** screen.

The screenshot shows the Cisco Cloud APIC interface. The navigation sidebar on the left is expanded to show the 'Contracts' category. The main content area displays the 'Configuration Drifts' tab, which includes a 'Detection Summary' section and a table of contracts with drifts.

**Detection Summary**

|                   |   |                      |   |
|-------------------|---|----------------------|---|
| Contracts Checked | 6 | Contracts With Drift | 5 |
|-------------------|---|----------------------|---|

**Filter by attributes**

| Status    | Contract                     |
|-----------|------------------------------|
| Transient | c_1<br>tn1                   |
| Raised    | ssh-http-https-icmp<br>infra |
| Raised    | netconf-ssh<br>infra         |

- In the **Navigation** sidebar, expand the **Application Management** category.
- From the **Application Management** category, select **Contracts**.
- In the **Contracts** screen, select the **Configuration Drifts** tab.

In the **Configuration Drifts** tab, you can see a summary of any configuration issues with the contracts in your fabric.

For each contract with a drift, you will see the number of missing configurations and the severity of the issue.

You can refresh the information by clicking the refresh button in the top right of the main window.

**Step 3** In the **Configuration Drifts** screen, click the name of a contract to view its details, including the configuration drift issues.

**Step 4** In the **Contract details** view that opens, select the **Cloud Mapping** tab.

The **Cloud Mapping** view displays all the information about the contract and the cloud resources it uses.

**Contract ssh-http-https-icmp**

Overview Topology Cloud Resources Application Management **Cloud Mapping** BETA Event Analytics

Detection of configuration drifts is still in beta.

**Detection Summary**

|                            |                            |                          |                               |
|----------------------------|----------------------------|--------------------------|-------------------------------|
| Configuration Drift Status | Configured Cloud Resources | Expected Cloud Resources | Last Cloud Inventory Update   |
| 32 Drifts Found            | -                          | 32                       | Jun 23 2020 04:15:31pm -07:00 |

**Configuration Drifts**

| Status | Resource Type | Protocol | Port Range  | Source    | Destination                                               | Consumer EPG                     | Provider EPG                      | Drift Type          | Description                                 | Recommendati...                                      |
|--------|---------------|----------|-------------|-----------|-----------------------------------------------------------|----------------------------------|-----------------------------------|---------------------|---------------------------------------------|------------------------------------------------------|
| Raised | Inbound Rule  | TCP      | http        | 0.0.0.0/0 | uni/tn-infra/clo-udapp-cloud-infra/cloudepg-infra-routers | ext-networks infra > cloud-infra | infra-routers infra > cloud-infra | Deployment mismatch | Inbound rule missing at cloud provider site | Repost the configuration associated to this contract |
| Raised | Inbound Rule  | TCP      | ssh         | 0.0.0.0/0 | uni/tn-infra/clo-udapp-cloud-infra/cloudepg-infra-routers | ext-networks infra > cloud-infra | infra-routers infra > cloud-infra | Deployment mismatch | Inbound rule missing at cloud provider site | Repost the configuration associated to this contract |
| Raised | Inbound Rule  | ICMP     | unspecified | 0.0.0.0/0 | uni/tn-infra/clo-udapp-cloud-infra/cloudepg-infra-routers | ext-networks infra > cloud-infra | infra-routers infra > cloud-infra | Deployment mismatch | Inbound rule missing at cloud provider site | Repost the configuration associated to this contract |

The screen is divided into three sections, **Detection Summary**, **Configuration Drifts**, and **Mapped Cloud Resources**. Each section contains a table that lists the respective information about the contract you selected.

The **Detection Summary** table 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 **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 **Mapped 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.

# 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"}
```





# APPENDIX A

## Cisco Cloud APIC Error Codes

- [Cisco Cloud APIC Error Codes](#), on page 165

### Cisco Cloud APIC Error Codes

This section describes the Cisco Cloud APIC error codes.

**Table 36: Cisco Cloud APIC Error Codes**

| Component      | Error Code                                  | Constraint                                                                                                                            |
|----------------|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| cloud-template | CT_INFRANETWORK_COUNT                       | The count of <code>cloudtemplateInfraNetwork</code> MO is at most 1                                                                   |
| cloud-template | CT_INFRANETWORK_VRF                         | In the <code>cloudtemplateInfraNetwork</code> MO, the <code>vrfName</code> must be <code>overlay-1</code>                             |
| cloud-template | CT_INFRANETWORK_PARENT                      | For the <code>cloudtemplateInfraNetworkMO</code> , the parent MO must be <code>uni/tn-infra</code>                                    |
| cloud-template | CT_INFRANETWORK_NUMROUTERSPERREGION_MINIMUM | In the <code>cloudtemplateInfraNetwork</code> MO, for the attribute <code>numRoutersPerRegion</code> , the minimum allowed value is 2 |
| 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_INTNETWORK_COUNT                         | The count of <code>cloudtemplateIntNetwork</code> MO is at most 1                                                                     |

| Component      | Error Code                           | Constraint                                                                                                                                                                                                            |
|----------------|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cloud-template | CT_EXTNETWORK_COUNT                  | The count of <code>cloudtemplateExtNetwork MO</code> is at most 1                                                                                                                                                     |
| cloud-template | CT_VPNNETWORK_COUNT                  | The count of <code>cloudtemplateVpnNetwork MO</code> is at most 1                                                                                                                                                     |
| cloud-template | CT_OSPF_COUNT                        | The count of <code>cloudtemplateOspf MO</code> 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>         |
| 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 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_EXTNETWORK_REQUIRES_EXTSUBNETPOOL | The presence of <code>cloudtemplateExtNetwork</code> requires presence of <code>cloudtemplateExtSubnetPool</code>                                                                                                     |
| cloud-template | CT_EXTSUBNETPOOL_COUNT               | The count of <code>cloudtemplateExtSubnetPool</code> is at most 1                                                                                                                                                     |
| cloud-template | CT_EXTSUBNETPOOL_SUBNETPOOL_ADDRESS  | In <code>cloudtemplateExtSubnetPool</code> , the <code>subnetpool</code> must contain a network address.                                                                                                              |

| Component      | Error Code                               | Constraint                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cloud-template | CT_EXTSUBNETPOOL_SUBNETPOOL_IP_VERSION   | In <code>cloudtemplateExtSubnetPool</code> , the subnetpool must contain an 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 (i.e. the netmask must be 22 or less).                                                                                                                                                                                                                                                                                                  |
| cloud-template | CT_EXTSUBNETPOOL_AND_REMOTESITE          | The <code>cloudtemplateExtSubnetPool</code> needs to be big enough to have at least one <code>cloudtemplateRemoteSiteSubnetPool</code> for each <code>cloudtemplateRemoteSite</code> .                                                                                                                                                                                                                                   |
| cloud-template | CT_INTNETWORK_MISSING_HOME               | If there are any <code>cloudRegionName</code> under <code>cloudtemplateIntNetwork</code> , then one of the <code>cloudRegionName</code> must be associated to a region that is the home region of the cAPIC ( <code>capicDeployed</code> ).                                                                                                                                                                              |
| 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.                                                                                                                                                                                                                                                                                                                        |

| Component      | Error Code                                      | Constraint                                                                                                                                                                                                                                                                                                      |
|----------------|-------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 <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_AZURE_PROFILE_ROUTERUSERNAME_INVALID         | In Azure, some usernames are not valid (admin, root, ...) and must not end with a period.                                                                                                                                                                                                                       |
| cloud-template | CT_AZURE_PROFILE_ROUTERUSERNAME_TOO_LONG        | In Azure, username is restricted to max 20 characters.                                                                                                                                                                                                                                                          |
| 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>routerLicenseToken</code> must not contain invalid characters.                                                                                                                                                                                                 |
| 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-template | CT_PROFILE_ROUTERLICENSETOKEN_INVALID_CHARACTER | In <code>cloudtemplateProfile</code> , the <code>routerPassword</code> must be non-empty.                                                                                                                                                                                                                       |
| cloud-template | CT_APICSUBNET_INVALID_HOME_REGION               | In a <code>cloudApicSubnet</code> MO, the region marked for <code>capicDeployed</code> must be a valid region                                                                                                                                                                                                   |
| cloud-template | 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-template | CT_APICSUBNET_MULTIPLE_HOME_REGION      | In <code>cloudApicSubnet</code> MOs, at most 1 region may have <code>capicDeployed</code> as true.                                            |
| cloud-template | CT_HUBNETWORK_COUNT                     | The count of <code>cloudtemplateHubNetwork</code> MO is at most 1                                                                             |
| 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.                                      |
| cloud          | CLOUD_AZURE_CTXPROFILE_SUBNET_RENAME    | <code>cloudSubnet</code> name cannot be modified                                                                                              |
| cloud          | CLOUD_AZURE_CTXPROFILE_SUBNET_DUPLICATE | Two <code>cloudSubnet</code> in the same <code>cloudCtxProfile</code> cannot have the same name                                               |
| cloud          | CLOUD_CAPIC_IP_EXT_EPG_SELECTOR_MAXIMUM | There can be maximum 1 <code>cloudExtEpSelector</code> in the <code>cloudExtEPg</code> corresponding to Cloud APIC IP                         |
| cloud          | CLOUD_AZURE_ACCOUNT_IN_USE              | The association between the account and the tenant cannot be updated or deleted while the account is in use and context profiles are deployed |
| cloud          | CLOUD_AZURE_INFRA_ACCOUNT_CHANGE        | The account for the tenant infra cannot be modified or deleted                                                                                |
| cloud          | CLOUD_SOURCE_PORT_NOT_SUPPORTED         | Source port range is not allowed on Cloud APIC                                                                                                |
| cloud          | CLOUD_ONLY_PERMIT_ACTION_SUPPORTED      | Actions different from 'permit' are not supported on Cloud APIC                                                                               |

| Component      | Error Code                                      | Constraint                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------|-------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cloud          | CLOUD_CIDR_OVERLAP                              | The subnets of <code>cloudCidrs</code> cannot overlap                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| cloud          | CLOUD_SUBNET_USAGE                              | There can be at most 1 gateway subnet for a given zone and each user subnet should have exactly 1 gateway subnet in the same user subnet's zone                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| cloud          | CLOUD_AZURE_ACCOUNT_CRED_CROSS_TENANT           | The <code>cloudCredentials</code> used by the <code>cloudAccount</code> must be in the same tenant                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| cloud          | CLOUD_AZURE_ACCOUNT_AD_CROSS_TENANT             | The <code>cloudAd</code> used by the <code>cloudAccount</code> must be in the same tenant                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| cloud-template | CT_CLOUD_APICSUBNETPOOL_INSUFFICIENT_HUBNETWORK | 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. With <code>HubNetworking</code> enabled, there must be as many <code>cloudApicSubnetPool</code> as the <code>cloudRegionName</code> under <code>cloudtemplateIntNetwork</code> . |
| cloud          | CLOUD_SYSTEM_MO_IS_IMMUTABLE                    | Instances created by the system are immutable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| cloud-template | CT_BGPEVPN_PEERADDR_IP_VERSION                  | In <code>cloudtemplateBgpEvpn</code> , the <code>peeraddr</code> must contain a IPv4 address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| cloud-template | CT_BGPEVPN_PEERADDR_ADDRESS_TYPE                | In <code>cloudtemplateBgpEvpn</code> , the <code>peeraddr</code> IP address must be a host address                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| cloud          | CLOUD_APICSUBNETPOOL_SUBNET_HOST_PART           | In <code>cloudApicSubnetPool</code> <code>subnet</code> , the host part must be 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

| Component      | Error Code                                    | Constraint                                                                                                       |
|----------------|-----------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| cloud-template | CT_EXTSUBNETPOOL_CLOUD_APICSUBNETPOOL_OVERLAP | There is a subnet overlap between <code>cloudtemplateExtSubnetPool</code> and <code>cloudApicSubnetPool</code> . |

