Cisco Cloud Network Controller for Google Cloud User Guide, Release 25.0(5)

First Published: 2022-08-15

Last Modified: 2023-01-13

Trademarks

CONTENTS

New and Changed Information

This chapter contains the following sections:

New and Changed Information

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

Table 1: New Features and Changed Behavior in Cisco Cloud Network Controller for Release 25.0(5)

Feature or Change	Description	Where Documented
Product name change	Beginning with release 25.0(5), the Cisco Cloud APIC is renamed to Cisco Cloud Network Controller.	
Support for configuring a BGP-EVPN connection for inter-site connectivity	Beginning with release 25.0(5), for inter-site use cases, support is available for configuring a BGP-EVPN connection for inter-site connectivity between a Google Cloud site and other cloud sites or an ACI on-premises site, where Cisco Catalyst 8000Vs are used for the BGP-EVPN connection.	About Cisco Cloud Network Controller and Google Cloud, on page 5
Support for Cisco Catalyst 8000V router statistics	Beginning with release 25.0(5), you can view router statistics for individual cloud context profiles within a tenant. Statistics are displayed for Cisco Catalyst 8000V routers within the cloud context profile.	Cisco Cloud Network Controller Statistics, on page 123

About Cisco Cloud Network Controller

Overview

Cisco Application Centric Infrastructure (ACI) customers who own a private cloud sometimes may run part of their workload on a public cloud. However, migrating the workload to the public cloud requires working with a different interface and learning different ways to set up connectivity and define security policies. Meeting these challenges can result in increased operational cost and loss of consistency.

Cisco ACI can use Cisco Cloud Network Controller to extend a multi-site fabric to Amazon Web Services (AWS), Microsoft Azure, and Google Cloud public clouds.

What Cisco Cloud Network Controller Is

Cisco Cloud Network Controller is a software component of Cisco APIC that can be deployed on a cloud-based virtual machine (VM). Cisco Cloud Network Controller provides the following features:

- Provides an interface that is similar to the existing Cisco APIC to interact with the Google Cloud public cloud.

- Automates the deployment and configuration of cloud connectivity.

- Configures the cloud router control plane.

- Translates Cisco ACI policies to cloud native policies.

- Discovers endpoints.

Guidelines and Limitations

This section contains the guidelines and limitations for Cisco Cloud Network Controller.

- 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 Network

Controller virtual machines that managed the account. Cisco Cloud Network Controller can display stale cloud objects, but it cannot remove them. You must log in to the cloud account and remove them manually.

To check for stale cloud resources:

1. From the Cisco Cloud Network Controller 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, Topology, Cloud Resources, Application Management, and Event Analytics tabs appear.

3. Click the **Cloud Resources > Actions > View Stale Cloud Objects**. The **Stale Cloud Objects** dialog box appears.

About the Cisco Cloud Network Controller GUI

The Cisco Cloud Network Controller 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 **EPGs** 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 configuring Cisco Cloud Network Controller components, see

CHAPTER 3

About Cisco Cloud Network Controller and Google Cloud

The following topics in this chapter provide information on how Cisco Cloud Network Controller deployments work with Google Cloud.

Locating Important Google Cloud Project Information

After you create a Google Cloud project, that project will be assigned three unique identifiers:

- Project name
- Project ID
- Project number

You will need these three identifiers for your Google Cloud project at various points in the Cisco Cloud Network Controller configuration process. To locate the **Project Info** pane with these Google Cloud project identifiers, log into your Google Cloud account and select your particular Google Cloud project in the **Select a project** window. The **Dashboard** for this project is displayed, which provides the Project Info pane with these three unique identifiers for your Google Cloud project.
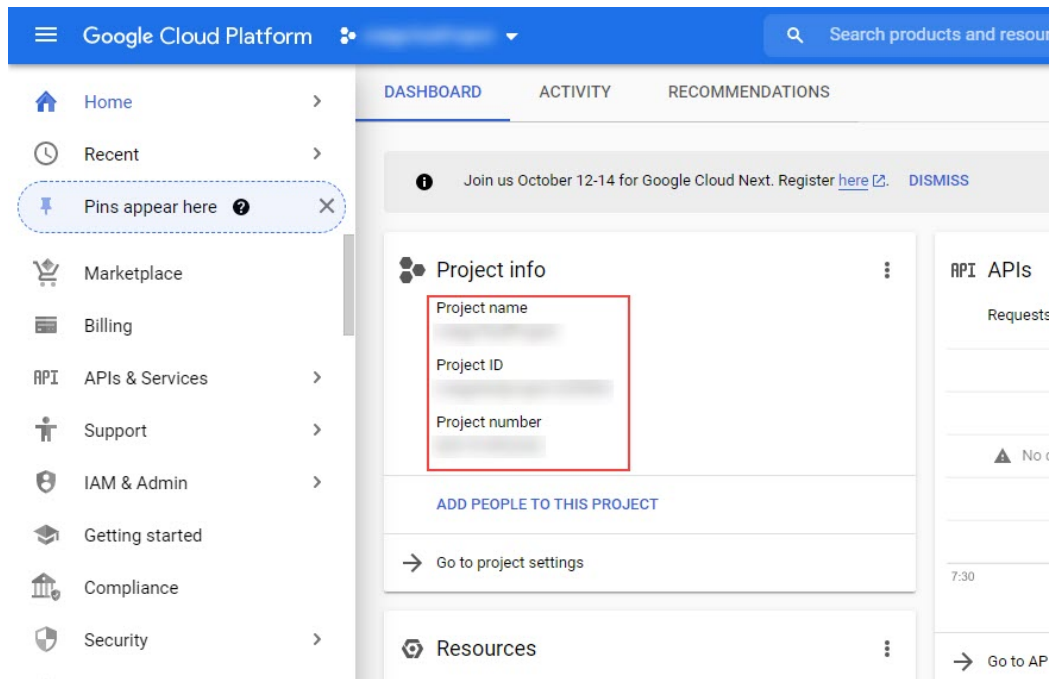
Understanding Google Cloud Deployments with Cisco Cloud Network Controller

Google Cloud organizes resources in a way that resembles a file system, where:

- The *Organization* at the top level can have multiple *Folders*.

- Every *Folder* can contain other *Folders*, or can contain *Projects*, where every *Project* has a unique ID.

- Cloud *resources* (such as VMs, VPCs, and subnets) are contained within a *Project*.

While the Organization and Folder levels are useful areas to understand from the Google Cloud perspective, the Project level is the most relevant from the Cisco Cloud Network Controller perspective.

Each Cisco Cloud Network Controller tenant is mapped one-to-one to a Google Cloud Project, which means that:

- A Cisco Cloud Network Controller tenant cannot span multiple Google Cloud Projects

- There cannot be more than one Cisco Cloud Network Controller tenant in a Google Cloud Project

With Cisco Cloud Network Controller, Google Cloud provides access to Projects using **Service Accounts**. These accounts are meant for applications that need to access Google Cloud services. They can be used to run and deploy Cisco Cloud Network Controller and to push policies for other tenants. Service accounts used in applications running within Google Cloud do not need credentials, whereas applications that are run external to Google Cloud need a pre-generated private key. Service Accounts reside in one Google Cloud Project, but they can also be given access to manage policies for other Projects (for Cisco Cloud Network Controller, other tenants).

The following sections provide more information on different ways that Cisco Cloud Network Controller tenants can be configured with Google Cloud:

User Tenants With Managed Credentials

This type of user tenant has the following characteristics:

- This tenant account is managed by the Cisco Cloud Network Controller.
- You will first choose **Managed Identity** in the Cisco Cloud Network Controller GUI as part of the tenant configuration process for this type of user tenant.
- After you have configured the necessary parameters in the Cisco Cloud Network Controller, you must then set the necessary roles for this tenant in Google Cloud. Add the service account created by the Cisco Cloud Network Controller as an IAM user with the following rules:
    - Cloud Functions Service Agent
    - Compute Instance Admin (v1)
    - Compute Network Admin
    - Compute Security Admin
    - Logging Admin
    - Pub/Sub Admin
    - Storage Admin

For instructions on creating this sort of tenant, see Creating a Managed Tenant Using the Cisco Cloud Network Controller GUI, on page 41.

User Tenants With Unmanaged Credentials

This type of user tenant has the following characteristics:

- This tenant account is not managed by the Cisco Cloud Network Controller.
- Before configuring the necessary parameters in the Cisco Cloud Network Controller for this type of tenant, you must first download the JSON file that contains the necessary private key information from Google Cloud for the service account associated with this tenant.
- You will then choose **Unmanaged Identity** in the Cisco Cloud Network Controller GUI as part of the tenant configuration process for this type of user tenant. As part of the configuration process for this type of tenant in Cisco Cloud Network Controller, you will provide the following information from the downloaded JSON file:
    - Key ID
    - RSA Private Key
    - Client ID
    - Email

For instructions on creating this sort of tenant, see Creating an Unmanaged Tenant Using the Cisco Cloud Network Controller GUI, on page 45.

External Network Connectivity Using Cloud Native Routers

Support is available for external connectivity between a Google Cloud site and non-Google Cloud sites or an external device. You can have this IPv4 connection by creating a VPN connection between a Google Cloud router and an external device.

The following sections provide more information on the components that allow for the new external network connectivity:

External VRF

An **external VRF** is a unique VRF that does not have any presence in the cloud. This VRF is not referred to in any cloud context profile used by Cisco Cloud Network Controller.

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

Cloud Native Routers

When configuring Cisco Cloud Network Controller with Google Cloud, the infra VPC uses Google Cloud native routers (Cloud Router and Cloud VPN gateway) to create IPsec tunnels and BGP sessions to on-premises sites, other cloud sites, or any remote device. Only IPv4 connectivity is supported for this type of connectivity using cloud native routers, where IPv4 sessions are created on an external VRF.

Google Cloud supports VPN connections both with static routes and with BGP. To create a VPN connection with BGP, Cisco Cloud Network Controller needs both a Cloud Router and a VPN gateway. A VPC can have multiple Cloud Routers and VPN gateways. However, Google Cloud has a restriction that both the Cloud Routers and the VPN gateways must be in the same region and in the same VPC. In addition, Cisco Cloud Network Controller has a restriction where only one cloud router and one cloud VPN gateway is supported per region.

VPN Communication

When configuring Cisco Cloud Network Controller with Google Cloud, the infra VPC is used to host the Cisco Cloud Network Controller and to host the VPN connections to external devices and sites. However, the infra VPC is not used as a transit to implement spoke-to-spoke communication. Instead, when configuring Cisco Cloud Network Controller with Google Cloud, spoke-to-spoke communication is done though spoke-to-spoke VPC peering.

The infra VPC uses the Google Cloud Router and Google Cloud VPN Gateway to create IPsec tunnels and BGP sessions to on-premises sites or to other cloud sites. Spoke VPCs peer with the infra VPC to share the VPN connections to external sites, where:

- Routes received on the VPN connections are leaked to the spoke VPCs

- Spoke VPC routes are advertised on the VPN connections

Using inter-VRF routing, the route is leaked between the external VRF of the VPN connections and the cloud local spoke VRFs.

A VPN gateway has two interfaces, and Google Cloud allocates public IP addresses to each of the interfaces. While the Google Cloud VPN gateway could have one or two interfaces, Cisco Cloud Network Controller only supports VPN gateways with two interfaces because two interfaces are required to achieve high availability.
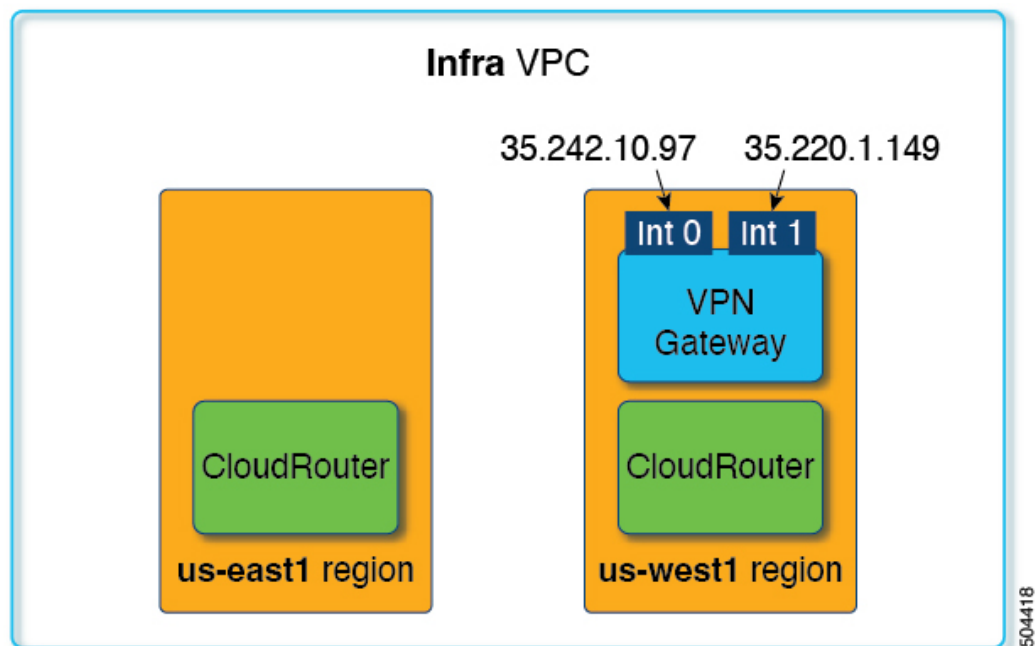
Hub Network Configuration

Rather than creating the hub network in a region based on the spoke attachments, the `cloudRegionName` MOs under a `cloudtemplateHubNetworkName` represents the regions where the hub network will be deployed, where `cloudtemplateHubNetworkName` represents a Google Cloud Router. Cisco Cloud Network Controller has a restriction of only one `cloudtemplateHubNetworkName`.

The hub network provides a way for establishing connectivity to an external site. Creating a hub network is a pre-requisite to creating an external network. You can create a hub network by specifying a name for the hub and the regions where the hub network should be deployed. For example, you may choose to deploy the hub network in us-central1 and us-east1. Cisco Cloud Network Controller will provision the Google Cloud Routers in these regions. Remember that only one hub network can be created, which means that Cisco Cloud Network Controller will only deploy one Cloud Router per region.

The following POST shows an example of network connectivity using this model. The `cloudtemplateHubNetwork` is used to create the hub network. In this example, the hub network is deployed

in four regions. External networks are created from each of the four regions using the
`cloudtemplateExtNetwork` MOs.

```
<polUni>
    <fvTenant name="infra" status="">
        <fvCtx name="extv1" pcEnfPref="enforced" status=""/>
         <fvCtx name="extv2" pcEnfPref="enforced" status=""/>
          <fvCtx name="extv3" pcEnfPref="enforced" status=""/>

          <cloudtemplateInfraNetwork name="default" vrfName="overlay-1"
hostRouterMode="manual" status="">
          <cloudtemplateIpSecTunnelSubnetPool subnetpool= "169.254.7.0/24" poolname="pool1"
 />
           <cloudtemplateIpSecTunnelSubnetPool subnetpool= "169.254.8.0/24" poolname="pool2"
 />
              <cloudtemplateIpSecTunnelSubnetPool subnetpool= "169.254.10.0/24"
poolname="pool3" />

              <cloudtemplateHubNetwork name="default" status="" >
                  <cloudtemplateHubNetworkName name="foo1" asn="64514" status="">
                      <cloudRegionName provider="gcp" region="us-west4" status="" />
                      <cloudRegionName provider="gcp" region="us-west2" status="" />
                      <cloudRegionName provider="gcp" region="us-east1" status="" />
                      <cloudRegionName provider="gcp" region="us-west1" status=""/>
                  </cloudtemplateHubNetworkName>
              </cloudtemplateHubNetwork>

              <cloudtemplateIntNetwork name="default">
                  <cloudRegionName provider="gcp" region="us-west1">
                      <cloudtemplateVpnRouter name="default" status=""/>
                  </cloudRegionName>
                  <cloudRegionName provider="gcp" region="us-west2">
                      <cloudtemplateVpnRouter name="default" status=""/>
                  </cloudRegionName>
                  <cloudRegionName provider="gcp" region="us-east1">
                      <cloudtemplateVpnRouter name="default" status=""/>
                   </cloudRegionName>
                  <cloudRegionName provider="gcp" region="us-west4">
                      <cloudtemplateVpnRouter name="default" status=""/>
                  </cloudRegionName>
              </cloudtemplateIntNetwork>

              <cloudtemplateExtNetwork name="default">
              </cloudtemplateExtNetwork>
               <cloudtemplateExtNetwork name="extnwfoo1" vrfName="extv1" hubNetworkName="foo1"
 vpnRouterName="default" status="">
                      <cloudRegionName provider="gcp" region="us-west1" status=""/>
                        <cloudtemplateVpnNetwork name="onprem01" remoteSiteId="1" status="">

                      <cloudtemplateIpSecTunnel peeraddr="128.1.1.1" preSharedKey="abcd"
poolname="pool1" status="">
                         <cloudtemplateBgpIpv4 peeraddr="0.0.0.0/0" peerasn="64529" status=""/>

                      </cloudtemplateIpSecTunnel>
                  </cloudtemplateVpnNetwork>
              </cloudtemplateExtNetwork>
              <cloudtemplateExtNetwork name="extnwfoo2" vrfName="extv2" hubNetworkName="foo1"
 vpnRouterName="default" status="">
               <cloudRegionName provider="gcp" region="us-west2" status=""/>
                  <cloudtemplateVpnNetwork name="onprem02" remoteSiteId="2" status="">
                        <cloudtemplateIpSecTunnel peeraddr="128.1.1.2" preSharedKey="def"
poolname="pool2" status="">
                      <cloudtemplateBgpIpv4 peeraddr="0.0.0.0/0" peerasn="64529" status=""/>
```

```
                    </cloudtemplateIpSecTunnel>
                </cloudtemplateVpnNetwork>
             </cloudtemplateExtNetwork>
            <cloudtemplateExtNetwork name="extnwfoo3" vrfName="extv3" hubNetworkName="foo1"
 vpnRouterName="default" status="">
                <cloudRegionName provider="gcp" region="us-east1" status=""/>
                   <cloudtemplateVpnNetwork name="onprem03" remoteSiteId="3" status="">
                       <cloudtemplateIpSecTunnel peeraddr="128.1.1.3" preSharedKey="abc"
poolname="pool3" status="">
                       <cloudtemplateBgpIpv4 peeraddr="0.0.0.0/0" peerasn="64529" status=""/>

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

In this example POST:

- **cloudtemplateExtNetwork**: You can have multiple `cloudtemplateExtNetwork` entries, each with a unique name, that represent an external network on an external VRF.

  Within the `cloudtemplateExtNetwork` area are the following fields:

  - **vrfName**: This property represents the VRF used for the external network (for example, a transport VRF). Multiple remote sites can use the same transport VRF, which means that all of these remotes sites are treated as one VRF on the cloud and all of the remote sites receive the same routes from the cloud.

  - **hubNetworkName**: This property represents the name of the hub network used by this external network. This name refers to one of the hub networks created in the `cloudtemplateHubNetworkName` area.

  - **vpnRouterName**: This property represents the name of the VPN router used by this external network. This name refers to the VPN router created by `cloudtemplateVpnRouter`.

  In addition, an external network can be deployed in multiple regions, and a router used on the external network should be deployed in those regions (in other words, `hubNetworkName` and `vpnRouterName` should exist in those regions).

- **cloudtemplateVpnNetwork**: This MO represents a remote site.

  Within the `cloudtemplateVpnNetwork` area is the **remoteSiteId** field. This property represents the remote site ID.

- **cloudtemplateVpnRouter**: This MO translates to a Google Cloud VPN gateway. Only one `cloudtemplateVpnRouter` is allowed, with the name `default`.

- **cloudtemplateIpSecTunnel**: This MO represents a remote peer.

- **cloudtemplateBgpIpv4**: This MO represents a remote site IPv4 BGP peer.

  If the `peeraddr` entry under `cloudtemplateBgpIpv4` has the default address (0.0.0.0/0), then the remote BGP peer is assumed to be the inner address of the tunnel on the remote device.

Note that the model above supports the following:

- Both ikev1 and ikev2 to an external device.

• Multiple `cloudtemplateIpSecTunnelSubnetPool` subnet pools. The allowed IP ranges in the `cloudtemplateIpSecTunnelSubnetPool` subnet pools is dependent on the cloud provider and use case. For example, 169.254.0.0/16 or a lesser subnet of it is supported for Google Cloud VPN connections.

Inter-Site Connectivity Using BGP-EVPN

Beginning with release 25.0(5), for inter-site use cases, support is also available for configuring a BGP-EVPN connection for inter-site connectivity in these scenarios:

• Cloud site-to-cloud site:

  • Google Cloud site-to-Google Cloud site

  • Google Cloud site-to-AWS site

  • Google Cloud site-to-Azure site

• Google Cloud site-to-ACI on-premises site

In each of these scenarios, Cisco Catalyst 8000Vs are used for the BGP-EVPN connection.
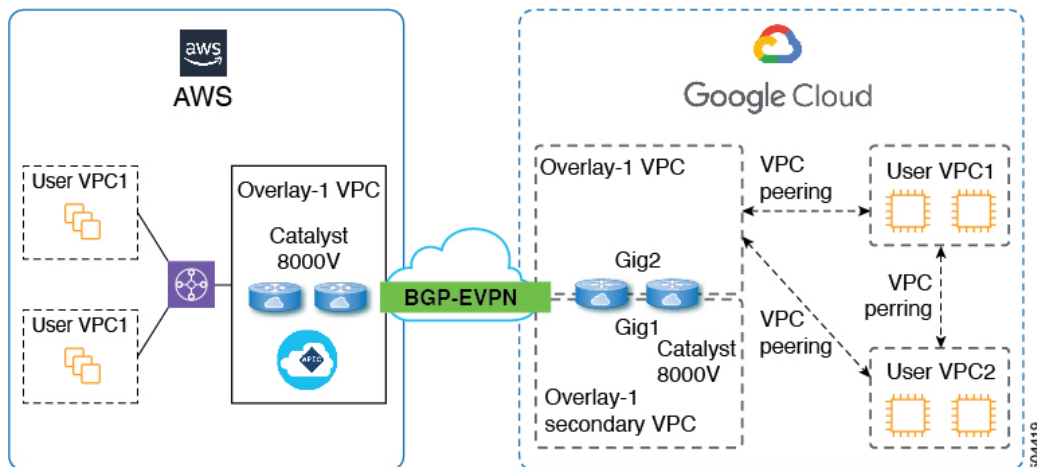
The following sections provide more information on the components that allow for inter-site connectivity using BGP-EVPN:

Characteristics of Inter-Site Connectivity Using BGP-EVPN

Based on Google Cloud behavior, each network interface of a VM or instance must be associated with a different VPC. Because the Cisco Catalyst 8000V is also a VM, this means that each network interface for a given Cisco Catalyst 8000V has to be associated with a different VPC. Two gigabit network interfaces in the Cisco Catalyst 8000V are therefore used in the following ways:

• The gig1 interface is associated with the overlay-1 secondary VPC. In addition, the gig1 interface is used as the management interface.

- The gig2 interface is associated with the overlay-1 VPC. In addition, the gig2 interface is used as the routing interface.

VPC Peering

In order to have communcation from the spoke VPC to an on-premises network, the spoke VPC must have peering enabled to the hub VPCs. The peering is automated by intent from Cisco Cloud Network Controller. VPC peering for Cisco Cloud Network Controller with Google Cloud employs a hub-spoke topology, as shown in the following figure.

Cisco Cloud Network Controller with Google Cloud uses three types of VPC peering:

- Spoke-to-spoke VPC peering: This is used for spoke-to-spoke intra-site communication.

- Hub-to-spoke VPC peering: This is used for inter-site communication that goes through the Cisco Catalyst 8000V routers using BGP-EVPN.

- Hub-to-hub VPC peering: This is used for communication between the Cisco Cloud Network Controller in the overlay-1 VPC and the Cisco Catalyst 8000V routers management interfaces in the overlay-1 secondary VPC.

Note that the overlay-1 secondary VPC is not involved in the data path for either spoke-to-spoke or inter-site traffic.
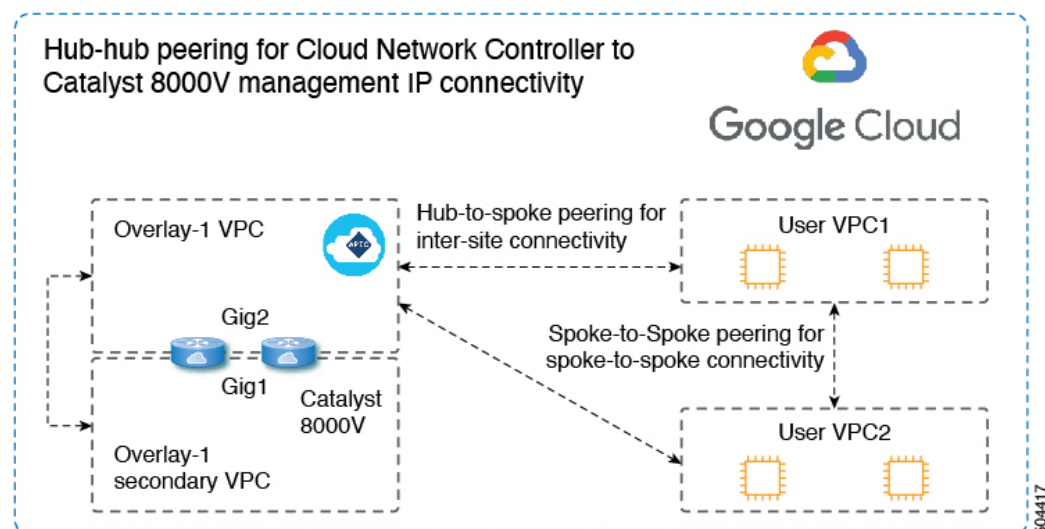
Cisco Cloud Network Controller automates configurations to exchange the routes between cloud sites in the following situations:

- Overlay-1 VPC to the destination in the same site: The overlay-1 VPC has the route to the spoke VPC in the same site through VPC peering.

- Spoke VPCs to the destination in another site: The routes for the subnets in the other site are added to the overlay-1 VPC by Cisco Cloud Network Controller and the routes are exported to the spoke VPCs. In this way, the spoke VPCs have the routes to reach the destination subnets in the other site.

- Between Cisco Catalyst 8000Vs in different sites: The static route for the spoke VPC CIDRs are added to the Cisco Catalyst 8000V routers in the same site. The static routes are redistributed to the Catalyst

8000V routers in the other site through BGP EVPN. In this way, the Catalyst 8000Vs have the routes to reach the destination subnets in the other site, as shown in the following figure.
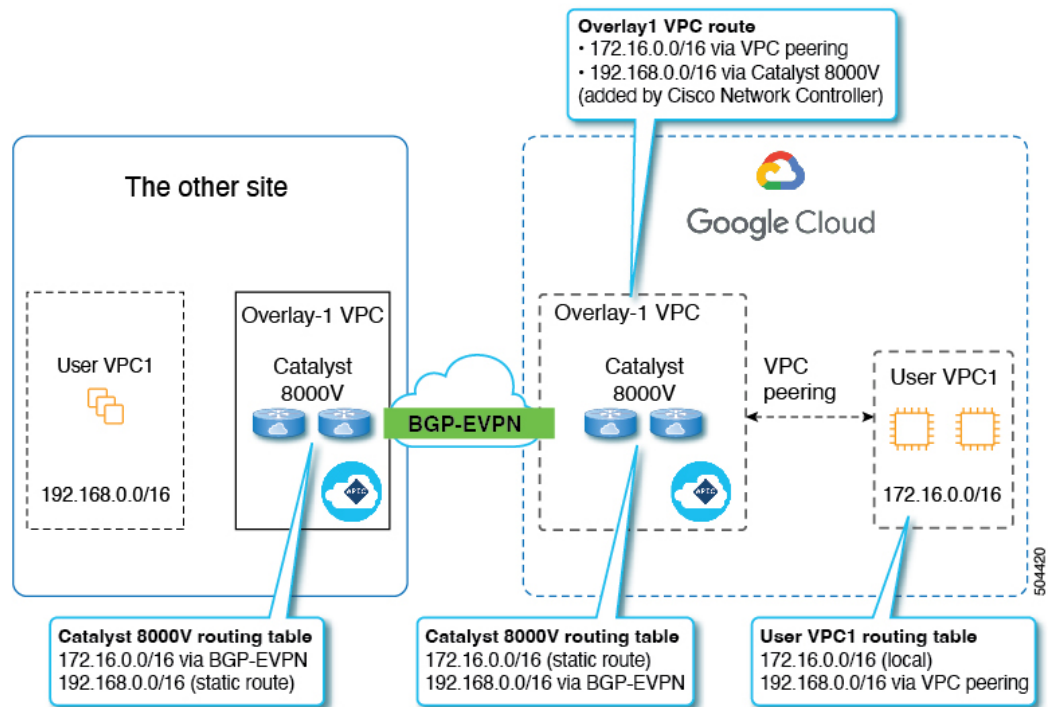
In this scenario, a static route to the remote CIDR is programmed in the hub VPC with the next hop as the Cisco Catalyst 8000V. These routes are learned by the spoke VPC using peering.

Configuring Routing and Security Policies Separately

To allow communication between two endpoints in different VRFs, you need to establish routing and security policies separately:

- **Routing policies**: Policies used to define routes to establish traffic flow

- **Security policies**: Rules used for security purposes, such as zoning rules, security-group rules, ACLs, and so on

For Google Cloud, routing must be configured independent of security. In other words, for Google Cloud, "contracts" are used only for security. To configure routing, you must configure route-maps.

Configuring Routing Policies

Using inter-VRF routing, you can configure an independent routing policy to specify which routes to leak between a pair of VRFs. To establish routing, you must configure route maps between a pair of VRFs.

For situations where you can use route maps to set which routes to leak between a pair of VRFs, the following types of VRFs are used for inter-VRF routing:

- **External VRF** is a VRF that is associated with one or more external networks.

- **Internal VRF** is a VRF that has one or more cloud context profiles or cloud subnets associated with it.

When configuring inter-VRF routing with these types of VRFs:

- Between a pair of internal VRFs, you must always leak all routes.

- From an internal VRF to an external VRF, you can leak specific routes or all routes.

- From an external VRF to an internal VRF, you must leak all routes.

Guidelines and Restrictions

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

- Routes are always leaked bi-directionally between two VRFs. For every route leak entry from one tenant/VRF under another tenant/VRF, there must be a corresponding route leak entry going in the opposite direction.

  For example, assume there are two tenants (`t1` and `t2`) and two corresponding VRFs (`v1` and `v2`). For every route leak entry `t1:v1` under the VRF `t2:v2`, there must be a corresponding route leak entry `t2:v2` under the VRF `t1:v1`.

- Once you associate an external VRF with an external network, if you want to change the external VRF, you need to delete the external network and then recreate the external network with the new external VRF.

- You cannot configure "smaller" prefixes to be leaked while a "larger" prefix is already being leaked. For example, configuring the 10.10.10.0/24 prefix will be rejected if you already have the 10.10.0.0/16 prefix configured to be leaked. Similarly, if you configure the 0.0.0.0/0 (leak all) prefix, no other prefix will be allowed to be configured.

Configuring Security Policies

While an EPG in Cisco Cloud Network Controller corresponds to security groups in AWS and Azure, there is no equivalent corresponding component in Google Cloud for an EPG. The closest equivalent in Google Cloud is a combination of firewall rules and network tags.

The firewall resource in Google Cloud is global to the project (tenant). Firewall rules are associated with a single VPC and their scope applies to the entire VPC globally. The scope of the firewall rule is further defined by the Target parameter. In other words, the set of instances that a rule is applied to can be selected by one or more of the following Target types:

- **Network tags**: Network tags are key strings that drive the VM's firewall and routing configuration on Google Cloud. Instances (for example, VMs) can be tagged with unique strings. Firewall rules are applied to all instances with equal tags. Multiple tag values act as a logical 'or' operator, where the firewall rule is applied as long as at least one tag matches.

- **All instances in the network**: The firewall rule applies to all instances in the VPC.

Firewall rules also identify the source and destination of the traffic. Depending on whether the rule is for ingress traffic (going to a VM) or egress traffic (leaving a VM), the source and destination fields accept different values. The following list provides more information on those values:

- **Ingress rules**:

- **Source**: Can be identified using:

  - Network tags

  - IP addresses

  - A combination of IP addresses and network tags with a logical 'or' operator

- **Destination**: The Target parameter identifies the destination instances

- **Egress rules**:

  - **Source**: The Target parameter identifies the source instances

  - **Destination**: Can be identified using only IP addresses (not network tags)

How Cisco Cloud Network Controller Implements Firewall Rules With Google Cloud

The following list describes how Cisco Cloud Network Controller implements firewall rules with Google Cloud :

- **Global resources**: VPCs and firewalls in Google Cloud are global resources, so Cisco Cloud Network Controller does not have to program firewall rules for endpoints that span multiple regions. The same firewall rules apply for any region where the endpoint resides.

- **Firewall egress rules and network tags**: Firewall egress rules do not support network tags as a destination field, so you must list individual IP addresses for endpoints.

- **Source tags in firewall ingress rules and alias IP ranges**: Firewall ingress rules do not include the alias IP ranges of VMs matching the network tags used in the source field.

- **Priority fields in firewall rules**: Google Cloud evaluates firewall rules following their priority values.

Given that Google Cloud firewall rules follow a priority list, Cisco Cloud Network Controller configures a pair of low-priority deny-all ingress and egress rules when the VPC is created. Afterwards, Cisco Cloud Network Controller configures rules that open traffic according to the EPG's contracts with higher priority. Therefore, if there is no explicit rule that allows certain traffic as a result of an EPG contract, the low-priority rule matches and the default behavior is deny-all.

Endpoints and Endpoint Selectors

On the Cisco Cloud Network Controller, a cloud EPG is a collection of endpoints that share the same security policy. Cloud EPGs can have endpoints in one or more subnets and are tied to a VRF.

The Cisco Cloud Network Controller has a feature called endpoint selector, which is used to assign an endpoint to a Cloud EPG. The endpoint selector is essentially a set of rules run against the cloud instances assigned to the Google Cloud VPC managed by Cisco ACI. Any endpoint selector rules that match endpoint instances will assign that endpoint to the Cloud EPG. The endpoint selector is similar to the attribute-based microsegmentation available in Cisco ACI.

Following are the types of endpoint selectors available for the two types of cloud EPGs:

- **Application EPGs**:

  - **IP**: Used to select by the IP address or subnet.

  - **Region**: Used to select by the region of the endpoint.

- **Custom**: Used to select by a custom tag or label. For example, if you added a Location tag in Google Cloud, you might create the custom tag Location in this field to match the Location tag that you added in Google Cloud earlier.

- **External EPGs**:

  Subnet: The subnet selector is a type of endpoint selector where the match expression uses the IP address of a subnet, so an entire subnet is assigned as being part of the EPG. Essentially, when you use the subnet selector as the endpoint selector, all of the endpoints within that subnet belongs to the associated EPG.

When using Cisco Cloud Network Controller endpoint selectors with Google Cloud, a network tag is applied that associates the EPG to the matching VM in Google Cloud. Once the network tag is configured in the VM, Google Cloud applies the firewall rules for the VM's traffic.

VMs on Google Cloud also support labels. Labels are key-value pairs that are meant to be an organizational tool. The custom endpoint selector in Cisco Cloud Network Controller recognizes the labels assigned to the VMs in Google Cloud.

Cisco Cloud Network Controller reserves a unique network tag string for each EPG. In Google Cloud, this value is used as the target field in the firewall rules created for the EPG. When a new VM matches an endpoint selector of the EPG, Cisco Cloud Network Controller appends this value to the existing VM's network tags. In addition, the EPG's network tag is used in the source field of the Google Cloud firewall rules.

For example, consider the sample configuration below:

```
<cloudEPg  name="epg1" >
    <cloudRsCloudEPgCtx tnFvCtxName="v1"/>
    <fvRsProv tnVzBrCPName="httpSSHFamily"/>
    <cloudEPSelector name="web-selector" matchExpression="custom:server=='web'"/>
    <cloudEPSelector name="web-selector" matchExpression="custom:server==backend"/>
</cloudEPg>
<cloudEPg  name="epg2" status="">
    <cloudRsCloudEPgCtx tnFvCtxName="v1"/>
    <fvRsCons tnVzBrCPName="httpSSHFamily"/>
    <cloudEPSelector name="database-selector" matchExpression="custom:server=='database'"/>
</cloudEPg>
```

Assuming there are three endpoints in the VPC with the following configuration, Cisco Cloud Network Controller configures the following network tags, where the Cisco Cloud Network Controller-configured network tags are in the following format:

```
capic-<app-profile-name>-<epg-name>
```

Endpoint	Application Profile	EPG	Primary IP	Labels	Cisco Cloud Network Controller-Configured Network Tags
EP1	First application profile (app01)	First EPG (epg01)	10.0.0.1	server:web	capic-app01-epg01
EP2	Second application profile (app02)	Second EPG (epg02)	20.0.0.1	server:backend	capic-app02-epg02
EP3	Second application profile (app02)	Third EPG (epg03)	30.0.0.1	server:database	capic-app02-epg03

Cisco Cloud Network Controller needs admin permission over the VMs in order to set their network tags. This permission is granted by the *Compute Instance Admin* role.

There might be cases where Cisco Cloud Network Controller does not have this permission and cannot manage the VM's tags. In those scenarios, you can configure the network tags in your VMs first and then provide the proper endpoint selector configuration to Cisco Cloud Network Controller later on.

To see firewall rules:

- **In Google Cloud**: In your Google Cloud account, navigate to **VPC Network** > **Firewall**.

  - If the VM is part of an EPG, you can find the endpoints by expanding a firewall rule and then viewing the multiple entries shown in the **Filters** column. which are the endpoints.

  - Use the entry in the **Type** column to determine if a particular firewall rule is an ingress or an egress firewall rule.

    - If the firewall rule is an ingress type, then traffic is being sent to these endpoints.

    - If the firewall rule is an egress type, then these entries show where it can receive the traffic.

- **In Cisco Cloud Network Controller**: Firewall rules are associated with VPCs, so navigate to **Cloud Resources** > **VPCs**, then double-click on a VPC to get the detail screen. Then click on the **Cloud Resources** tab; there you will see the ingress and egress rules.

Understanding VPCs and Subnets Under Google Cloud and Cloud Context Profiles Under Cisco Cloud Network Controller

In Google Cloud, a VPC is a global resource, whereas a subnet is regional and spans every availability zone in the region, but a subnet cannot overlap with other subnets in the same VPC or in peered VPCs.

Each subnet must have exactly one primary CIDR block (IP range) and can have up to 30 secondary CIDR blocks. There can be up to 300 primary and secondary CIDRs in a VPC. The NIC for each VM gets its primary internal IP address from the primary CIDR block, whereas secondary IP ranges can only be used for **alias IP ranges**, which is a Google Cloud organizational tool to assign address pools to containers or applications running inside the VM.

The following provides more information on the associations between Cisco Cloud Network Controller objects and Google Cloud objects:

- **One-to-one mapping of Google Cloud VPC to Cisco Cloud Network Controller VRF**: A Google Cloud VPC is deployed for each Cisco Cloud Network Controller VRF (`fvCtx` object). Cloud context profiles (`cloudCtxProfile` object) define the set of regional subnets to deploy. Every cloud context profile in the same VRF maps to the same VPC.

- **Google Cloud subnets and their secondary IP ranges**: Cisco Cloud Network Controller deploys a subnet with primary and secondary IP ranges using Cisco Cloud Network Controller CIDR and subnet objects. The Cisco Cloud Network Controller subnet object is used to represent an IP range and the Cisco Cloud Network Controller CIDRs's primary property tells whether it is primary or secondary. Secondary Cisco Cloud Network Controller subnet objects are associated with the corresponding primary one, because only the latter deploys the actual subnet in Google Cloud.

Understanding VPC Groups

The cloud context profile is used within Cisco Cloud Network Controller as a mapping tool for a VPC, where one cloud context profile is associated with one VPC. The cloud context profile also contains information on the region association, where the cloud context profile is used to determine which region a VPC gets deployed to.

In Google Cloud, when you want to create a VPC, you might have to create multiple cloud context profiles through Cisco Cloud Network Controller if you want to deploy subnets in multiple regions. However, VPCs are global in nature with Google Cloud, where a VPC spans all the regions.

Therefore, a property called **VPC group** (`vpcGroup`) is available within the cloud context profile that allows Cisco Cloud Network Controller to group multiple cloud context profiles together to form one VPC. Multiple cloud context profiles that are associated with each other using the VPC group feature form the VPC construct within Google Cloud, where the VPC group name is shown in Google Cloud.

Since only one Google Cloud VPC is allowed within one Cisco Cloud Network Controller VRF, you must use the same name for the VPC group property for each cloud context profile listed in a VRF. Profiles having the same VPC group name reside in the same VPC.

The scope of this matching mechanism is at the tenant level. The same values can be reused across tenants, but they implicitly define different groups, since they are also part of different Google Cloud Projects.

Cisco Cloud Network Controller deploys a VPC for each distinct `fvCtx`, `cloudRsToCtx` and `vpcGroup` tuple, as long as there is at least one `cloudSubnet` defined. The cloud context profile becomes a container of regional resources, such as subnets, associated to a VRF, and it no longer maps to a VPC.

The example below defines two context profiles (c1 and c2) inside the same VRF (v1) with one VPC group (vpc-1). This configuration deploys one VPC, where the subnets defined in profiles c1 and c2 are deployed in that VPC because they are part of the same VPC group.

```
<fvTenant name="t1">
    <fvCtx name="v1"/>
    <cloudCtxProfile name="c1" vpcGroup="vpc-1">
        <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-gcp/region-us-west1" />
        <cloudRsToCtx tnFvCtxName="v1"/>
        <cloudCidr addr="10.0.0.0/16" primary="yes" >
            <cloudSubnet ip="10.0.1.0/24">
                <cloudRsZoneAttach
tDn="uni/clouddomp/provp-gcp/region-us-west1/zone-default"/>
            </cloudSubnet>
        </cloudCidr>
    </cloudCtxProfile>
    <cloudCtxProfile name="c2" vpcGroup="vpc-1">
        <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-gcp/region-us-east1" />
        <cloudRsToCtx tnFvCtxName="v1"/>
        <cloudCidr addr="20.0.0.0/16" primary="yes" >
            <cloudSubnet ip="20.0.1.0/24">
                <cloudRsZoneAttach
tDn="uni/clouddomp/provp-gcp/region-us-east1/zone-default"/>
            </cloudSubnet>
        </cloudCidr>
    </cloudCtxProfile>
</fvTenant>
```

Understanding Primary and Secondary Subnets and Subnet Groups

Cisco Cloud Network Controller deploys every subnet (`cloudSubnet`) in the VPC (which is identified by the tuple `fvCtx`, `cloudRsToCtx`, and `vpcGroup`) in the region that is pointed to by the `cloudRsCtxProfileToRegion` relation.

In Google Cloud, there is no concept of a primary CIDR for the VPC, but the **primary** flag in the CIDR (`cloudCidr`) field in the cloud context profile is available for Cisco Cloud Network Controller to support secondary IP ranges. Every subnet configured under a primary CIDR will be deployed as an actual Google Cloud subnet with the specified primary IP range (named *primary subnets*). Having multiple CIDRs set as primary under a given cloud context profile (`cloudCtxProfile`) is supported. Therefore, you can have more than one primary CIDR under a given cloud context profile with multiple primary subnets.

The following POST shows an example where one VPC and three subnets are deployed in Google Cloud.

```
<polUni>
    <fvTenant name="t1">
        <fvCtx name="v1"/>
        <cloudCtxProfile name="c1" vpcGroup="vpc-1">
            <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-gcp/region-us-west1" />
            <cloudRsToCtx tnFvCtxName="v1"/>
            <cloudCidr addr="10.0.0.0/16" primary="yes" >
                <cloudSubnet ip="10.0.1.0/24">
                    <cloudRsZoneAttach
tDn="uni/clouddomp/provp-gcp/region-us-west1/zone-default"/>
                </cloudSubnet>
                <cloudSubnet ip="10.0.2.0/24">
                    <cloudRsZoneAttach
tDn="uni/clouddomp/provp-gcp/region-us-west/zone-default"/>
                </cloudSubnet>
            </cloudCidr>
            <cloudCidr addr="20.0.0.0/16" primary="yes" >
                <cloudSubnet ip="20.0.1.0/24">
                    <cloudRsZoneAttach
tDn="uni/clouddomp/provp-gcp/region-us-west/zone-default"/>
                </cloudSubnet>
            </cloudCidr>
        </cloudCtxProfile>
    </fvTenant>
</polUni>
```

In the example above, one VPC is configured for the VRF `v1` with three primary subnets (10.0.1.0/24, 10.0.2.0/24, and 20.0.1.0/24) deployed in the us-west region.

A secondary CIDR contains the secondary IP ranges (called *secondary subnets*) that are configured in the existing primary subnets. When designating a CIDR as either primary or secondary, it's helpful to consider these differences between the two:

- The primary CIDR is normally the VM.

- The secondary CIDR is more of a container used for the application.

You can group together primary and secondary subnets into a **subnet group**. This grouping mechanism assigns secondary subnets (for example, IP ranges) to a primary subnet, which is mapped to an actual Google Cloud subnet. The scope of the subnet group is at the cloud context profile level. While you can have multiple cloud context profiles within the same tenant, subnets are part of a subnet group only within the same cloud context profile.

You will use the **subnet group label** to assign a unique label to a specific subnet group. If you have multiple subnets that have the same subnet group label, then those subnets all belong to the same subnet group as long

as they are all within the same cloud context profile. Note that while the subnet group label is used within Cisco Cloud Network Controller to group primary and secondary subnets, it is not used in Google Cloud.

Note the following guidelines for the primary and secondary CIDRs:

- **Primary CIDR**:

    - Any subnet group can have at maximum of only one subnet from the primary CIDR.

    - You can have multiple subnets in the primary CIDR, but all of the subnets must be in a separate subnet group.

- **Secondary CIDR**: You can have multiple subnets from the secondary CIDR in the same subnet group.

The following POST shows an example where two VPCs with two subnets each in different regions and having secondary CIDRs are deployed in Google Cloud.

```xml
<polUni>
    <fvTenant name="t1">
        <fvCtx name="v1"/>
        <fvCtx name="v2"/>
        <cloudCtxProfile name="c1" vpcGroup="vpc-1">
            <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-gcp/region-us-west1" />
            <cloudRsToCtx tnFvCtxName="v1"/>
            <cloudCidr addr="10.0.0.0/16" primary="yes" >
                <cloudSubnet ip="10.0.1.0/24" subnetGroup="subnet-1">
                    <cloudRsZoneAttach
tDn="uni/clouddomp/provp-gcp/region-us-west1/zone-default"/>
                </cloudSubnet>
                <cloudSubnet ip="10.0.2.0/24" subnetGroup="subnet-2">
                    <cloudRsZoneAttach
tDn="uni/clouddomp/provp-gcp/region-us-west1/zone-default"/>
                </cloudSubnet>
            </cloudCidr>
            <cloudCidr addr="40.0.0.0/16" primary="no">
                <cloudSubnet ip="40.0.1.0/24" subnetGroup="subnet-1">
                    <cloudRsZoneAttach
tDn="uni/clouddomp/provp-gcp/region-us-west1/zone-default"/>
                </cloudSubnet>
            </cloudCidr>
        </cloudCtxProfile>
        <cloudCtxProfile name="c2" vpcGroup="vpc-2">
            <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-gcp/region-us-east1" />
            <cloudRsToCtx tnFvCtxName="v2"/>
            <cloudCidr addr="20.0.0.0/16" primary="yes">
                <cloudSubnet ip="20.0.1.0/24" subnetGroup="subnet-1">
                    <cloudRsZoneAttach
tDn="uni/clouddomp/provp-gcp/region-us-east1/zone-default"/>
                </cloudSubnet>
            </cloudCidr>
            <cloudCidr addr="30.0.0.0/16" primary="no">
                <cloudSubnet ip="30.0.1.0/24" subnetGroup="subnet-1">
                    <cloudRsZoneAttach
tDn="uni/clouddomp/provp-gcp/region-us-east1/zone-default"/>
                </cloudSubnet>
            </cloudCidr>
        </cloudCtxProfile>
    </fvTenant>
</polUni>
```

Note that the subnet group `subnet-1` in the cloud context profile `c2` is not the same subnet group in the cloud context profile `c1`, because the scope of the subnet group is at the cloud context profile level.

The intent of the example above is summarized as follows:

- Tenant `t1` defines VRF `v1` and `v2`.

- Cloud context profile `c1` defines the subnets in region us-west1 for VRF `v1` and VPC group `vpc-1`. This deploys VPC `vpc-1`.

- Cloud context profile `c2` defines the subnets in region us-east1 for VRF `v2` and VPC group `vpc-2`. This deploys VPC `vpc-2`.

- The following subnets are deployed in VPC `vpc-1` in region us-west1:

    - `Subnet-1` subnet group:

        - Primary IP range: 10.0.1.0/24

        - Secondary IP ranges: 40.0.1.0/24

    - `Subnet-2` subnet group:

        - Primary IP range: 10.0.2.0/24

- The following subnets are deployed in VPC `vpc-2` in region us-east1:

    - Subnet-1:

        - Primary IP range: 20.0.1.0/24

        - Secondary IP ranges: 30.0.1.0/24

Guidelines and Limitations For Configuring Cisco Cloud Network Controller with Google Cloud

Following are the guidelines and limitations when configuring Cisco Cloud Network Controller with Google Cloud:

- For releases prior to release 25.0(5), Google Cloud does not support routing based on contracts. For more information, see Inter-Site Connectivity Using BGP-EVPN, on page 12.

- External connectivity between two Google Cloud sites is not supported.

- The external VRF can be configured only in the infra tenant in Cisco Cloud Network Controller.

- The tenant `common` in Cisco Cloud Network Controller cannot be associated with any Google Cloud project.

- In Google Cloud, the infra VPC and spoke VPCs are connected through VPC peering.

- In order to configure connectivity between the on-premises data center and the public cloud, you must manually configure the remote device by downloading the external device configuration files and manually enabling connectivity between Google Cloud and the external devices.

The external device configuration files that you download are not final configurations. Instead, the external device configuration files are provided more as a guidance. You must manually modify the information in the configuration files to configure the Google Cloud Router with IPSec, which is used to create connectivity between the on-premises data center and the public cloud, where:

- The Google Cloud Router and tunnels are deployed in the infra (hub) VPC.

- One cloud router per region is supported. Cloud routers can be deployed in a maximum of four regions.

- Spoke VPCs peer with the infra VPC to share the VPN connections to external sites, such as the on-premises data center.

Naming Length Restrictions Imposed By Google Cloud Firewall Rules

Google Cloud firewall rules are named resources, and Cisco Cloud Network Controller derives a name from the internal policy and uses that to deploy the Google Cloud firewall rules. Cisco Cloud Network Controller uses the following naming scheme for the internal policy:

```
{VPC-name}-{in/eg}-{target App-name}-{target EPG-name}-{contract-name}
```

The maximum length for a Google Cloud firewall rule name is 62 characters. This imposes a restriction on the names that you can use when configuring the following Cisco Cloud Network Controller components whose names are used in the Google Cloud firewall rule name:

- VPC group

- Application profile

- Application EPG or external EPG

- Contract

Knowing that the maximum number of characters is 62 for a Google Cloud firewall rule name, and taking into account the fixed areas in the string that makes up the Google Cloud firewall rule name:

- Hyphens (4 characters total)

- `in` (ingress) or `eg` (egress) value (2 characters)

That means that the total number of characters available for the combined names of all of the individual Cisco Cloud Network Controller components cannot exceed 56:

62 - 4 (number of hypens) - 2 (`in` or `eg` characters) = 56 characters

So, the sum of the lengths of the names of the VPC group, application profile, application EPG or external EPG, and contract must be smaller than 56 characters. On average, this allows for roughly 14 characters for the name of each component.

Cisco Cloud Network Controller Policy Model

About the ACI Policy Model

The ACI policy model enables the specification of application requirements policies. The Cisco Cloud Network Controller 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 Network Controller 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 Network Controller 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 Network Controller 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 1: ACI Policy Model Logical Constructs Overview
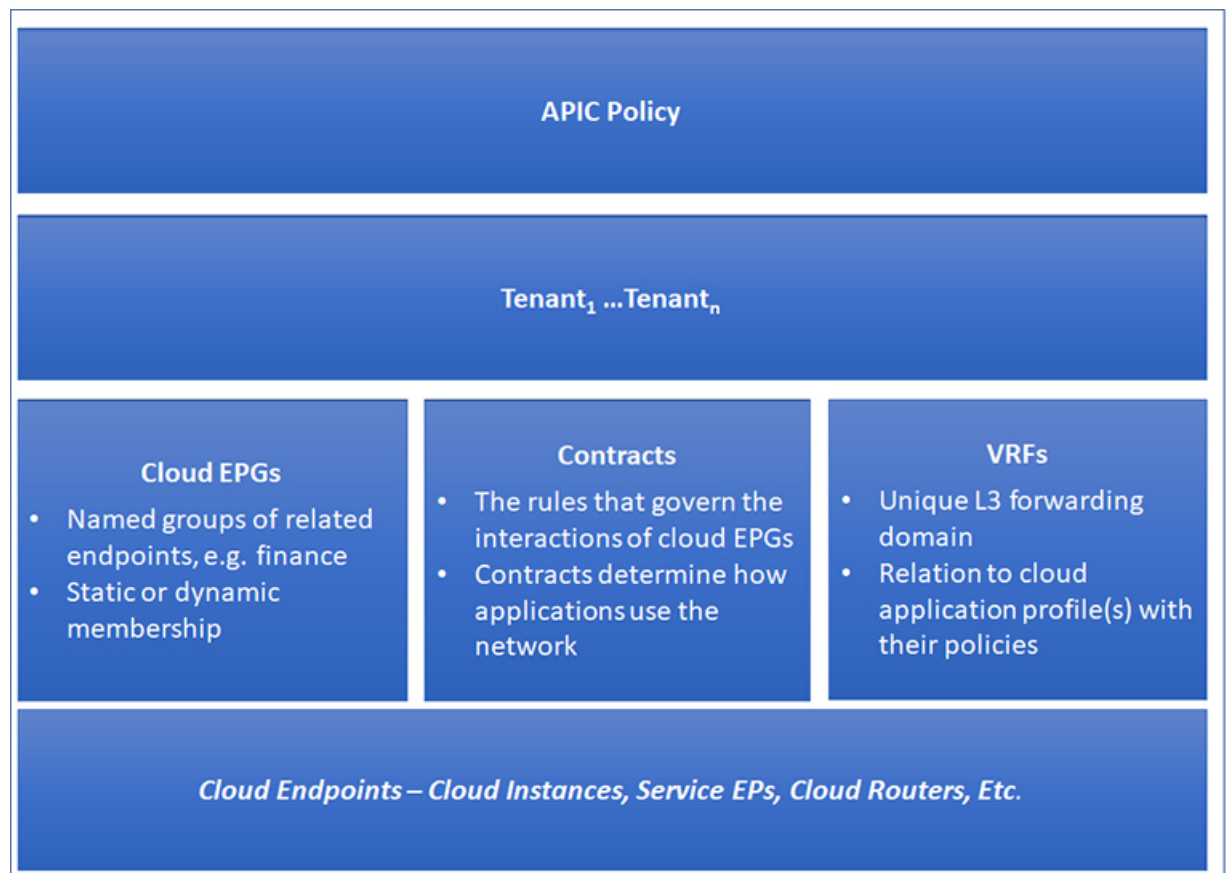
cloud infrastructure-wide or tenant administrators create predefined policies that contain application or shared resource requirements. These policies automate the provisioning of applications, network-attached services, security policies, and tenant subnets, which puts administrators in the position 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 Network Controller 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 Network Controller 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 2: Cisco ACI Policy Management Information Model Overview
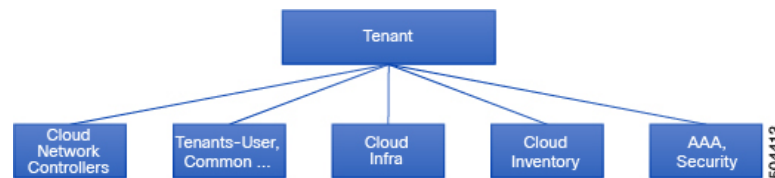
The hierarchical structure starts with the policy universe at the top (Root) and contains parent and child nodes. Each node in the tree is an MO and each object in the cloud infrastructure has a unique distinguished name (DN) that describes the object and locates its place in the tree.

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, intrusion detection appliances, and so on.

    - 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 Network Controller. For more information, see the *Cisco Cloud Network Controller 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.

- 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 Network Controller Security, on page 133

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 3: Tenants

Legend:
- Each box is an object
- Solid lines indicate that objects contain the objects below.
- Dotted lines indicate a relationship.
- 1:n indicates one to many; n:n indicates many to many.
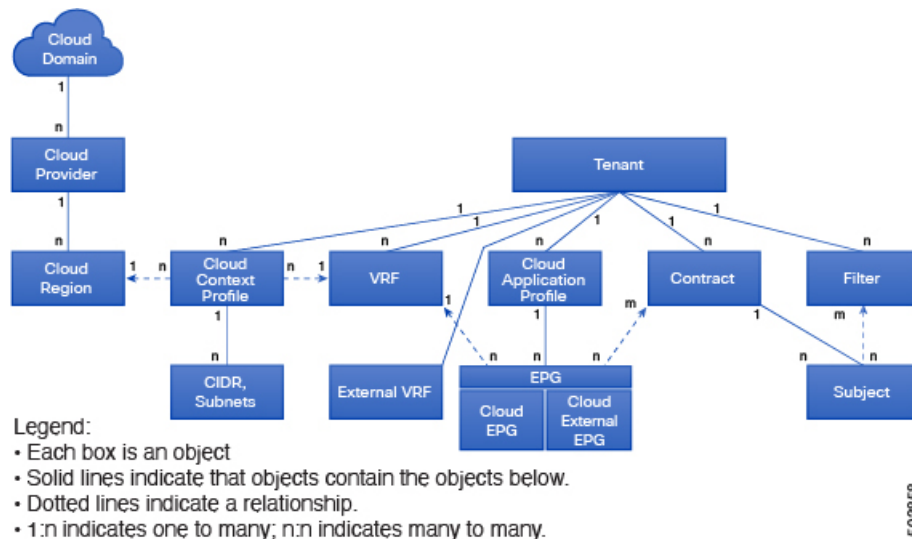
Tenants can be isolated from one another or can share resources. The primary elements that the tenant contains are filters, contracts, Virtual Routing and Forwarding (VRF) instances, cloud context profiles, Google Cloud provider configurations, and cloud application profiles that contain cloud endpoint 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 Google Cloud. A VPC 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. The ACI cloud infrastructure supports IPv4 and dual-stack configurations for tenant networking.

Cloud Context Profile

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

- CIDRs
- VRFs
- EPGs
- Regions
- VPCs
- Endpoints

VRFs

A Virtual Routing and Forwarding (VRF) object (`fvCtx`) or context is a tenant network (called a VRF in the Cisco Cloud Network Controller 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 4: VRFs
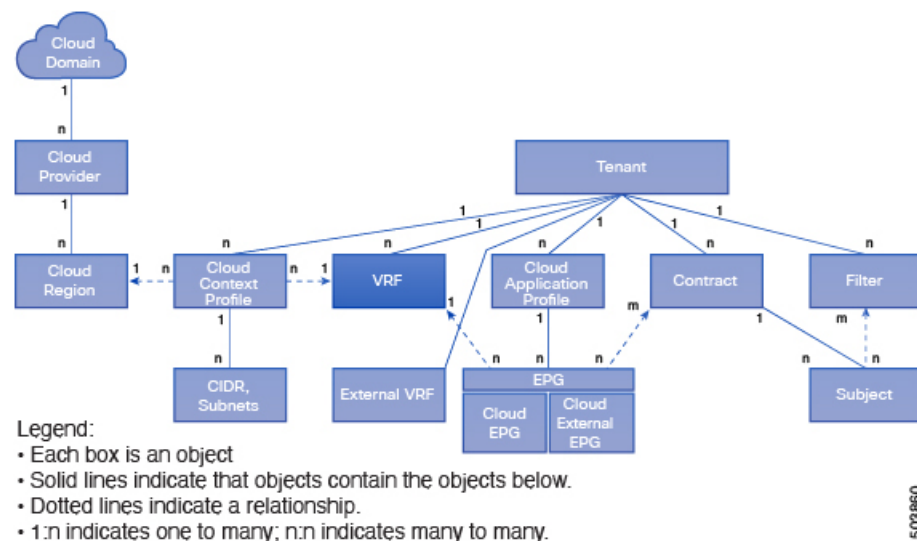
A VRF defines a Layer 3 address domain. One or more cloud context profiles are associated with a VRF. You can only associate one cloud context profile with a VRF in a given region. All the endpoints within the Layer 3 domain must have unique IP addresses because it is possible to forward packets directly 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.

External VRF

An **external VRF** is a type of VRF available for Cisco Cloud Network Controller. An external VRF is a unique VRF that does not have any presence in the cloud. This VRF is not referred to in any cloud context profile used by Cisco Cloud Network Controller.

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

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 5: Cloud Application Profiles
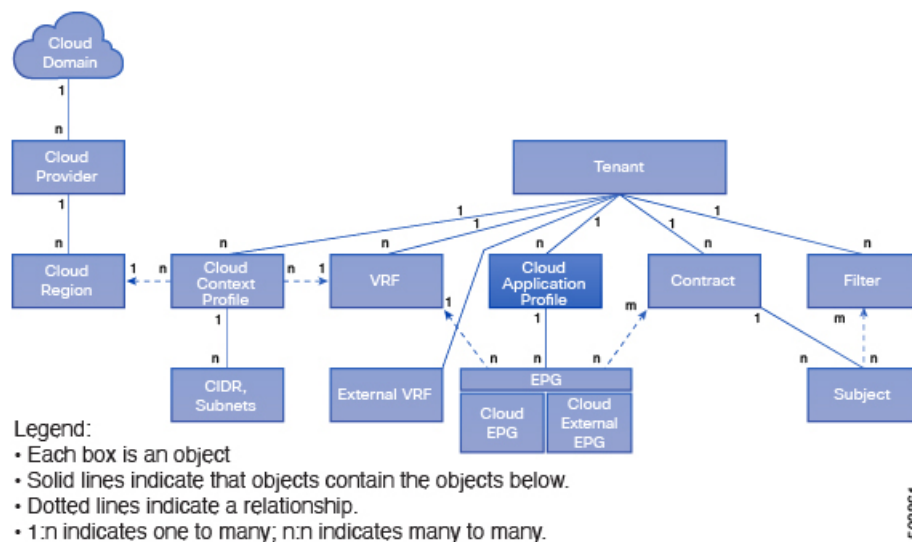
Cloud application profiles contain one or more cloud EPGs. Modern applications contain multiple components. For example, an e-commerce application could require a web server, a database server, data located in a storage service, and access to outside resources that enable financial transactions. The 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 6: Cloud Endpoint Groups
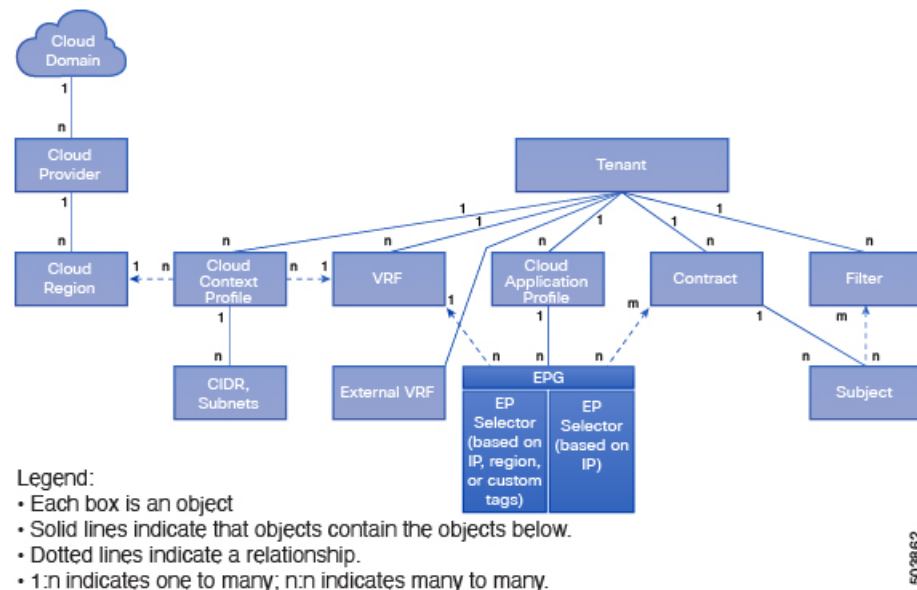
A cloud EPG is a managed object that is a named logical entity that contains a collection of endpoints. Endpoints are devices that are connected to the network. They have an address (identity), a location, attributes (such as version or patch level), and are virtual. Knowing the address of an endpoint 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 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 Network Controller 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 Google Cloud VPC managed by Cisco ACI. Any endpoint selector rules that match endpoint instances assign that endpoint to the Cloud EPG. The endpoint selector is similar to the attribute-based microsegmentation available in Cisco ACI.

Contracts

In addition to cloud EPGs, contracts (`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.
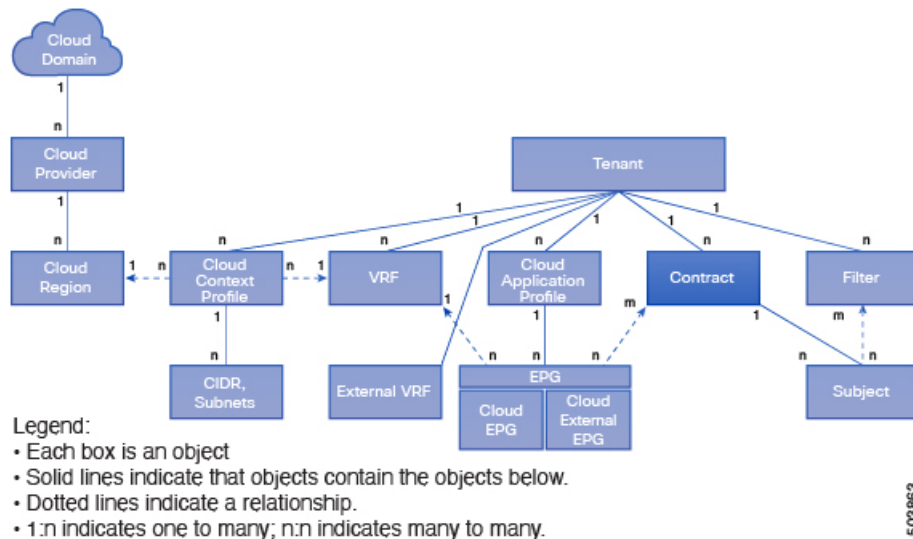
Figure 7: Contracts

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.

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 8: Subjects and Filters

Legend:
• Each box is an object
• Solid lines indicate that objects contain the objects below.
• Dotted lines indicate a relationship.
• 1:n indicates one to many; n:n indicates many to many.
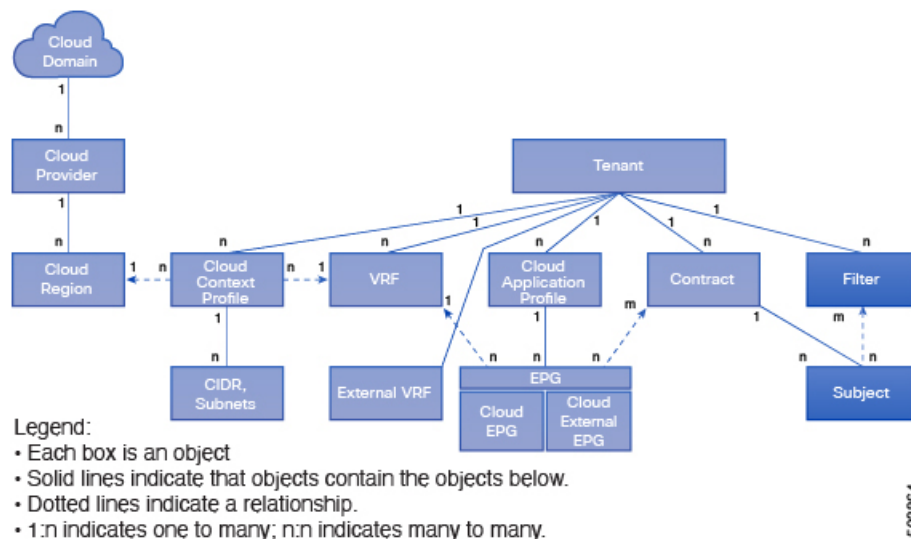
Contracts can contain multiple communication rules and multiple cloud EPGs can both consume and provide multiple contracts. A policy designer can compactly represent complex communication policies and re-use these policies across multiple instances of an application.

Note　Subjects are hidden in Cisco Cloud Network Controller and not configurable. For rules installed in Google Cloud, 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 Google Cloud constructs are always stateful, allowing return traffic.

About the Cloud Template

The cloud template provides a template that configures and manages the Cisco Cloud Network Controller 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 Network Controller 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 Google Cloud network configuration is the Virtual Private Cloud (VPC). Google Cloud supports many regions worldwide and one VPC is specific to one region.

The cloud template accepts one or more region names and generates the entire configuration for the infra VPCs in those regions. They are the infra VPCs. The Cisco Cloud Network Controller-managed object (MO) corresponding to the Google Cloud VPC is `cloudCtxProfile`. For every region specified in the cloud template, it generates the `cloudCtxProfile` configuration. A `cloudCtxProfile` is the topmost MO for all the configuration corresponding to a region. Underneath, it has many of other MOs organized as a tree to capture a specific configuration. The `cloudCtxProfile` MO for the infra VPC 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 a Google Cloud VPC is the CIDR. In Cisco Cloud Network Controller, you can choose and deploy CIDRs in the user VPCs. The CIDRs for the infra VPC are provided by users to the cloud template during the initial setup of the cloud site, and are deployed to the Google Cloud by the cloud template.

A property called `createdBy` is also available 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.

Multiple CIDR and subnet blocks can be configured on the infra VPC. You can create CIDRs and associate subnets in the infra VPC. The cloud template subnets will be mapped to the overlay-1 VRF. All subnets in the respective VRFs will have separate route tables in the cloud for VRF segregation.

For more information, see Creating an Application EPG Using the Cisco Cloud Network Controller GUI, on page 63.

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 2: Cloud Template MOs

MO	Purpose
`cloudtemplateInfraNetwork`	The root of the cloud template configuration. Attributes include: `numRoutersPerRegion`—The number of cloud routers for each `cloudRegionName` specified under `cloudtemplateIntNetwork`.
`cloudtemplateIntNetwork`	Contains a list of regions, which specify where you deploy the cloud routers. Each region is captured through a `cloudRegionName` child MO
`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 `cloudRegionName` child MO
`cloudtemplateIpSecTunnel`	Captures the IP address of the IPSec peer in the ACI on-premises site.

In Cisco Cloud Network Controller, 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 Network Controller, 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 9: Cloud and Cloud Template MO Conversion
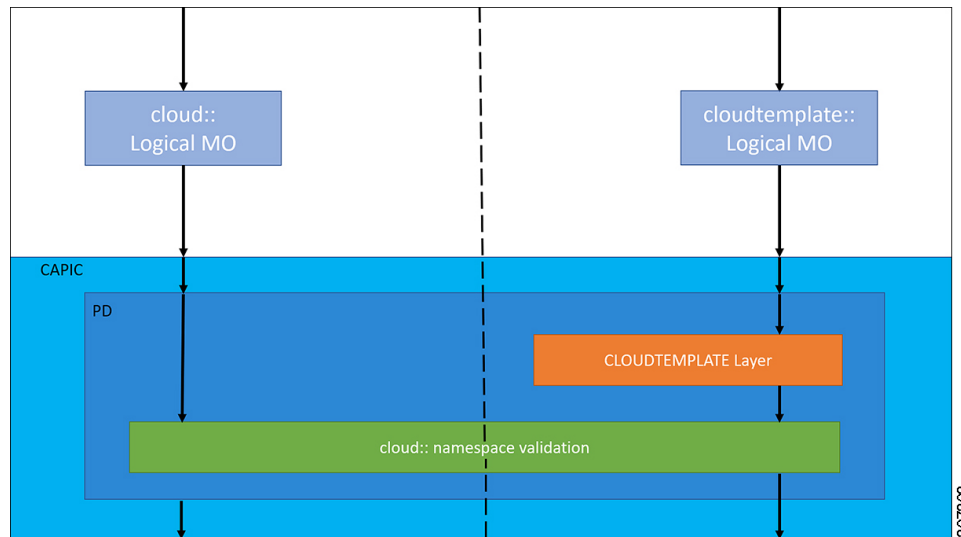
Note For information about configuring the cloud template, see Configuring Cisco Cloud Network Controller Components, on page 39

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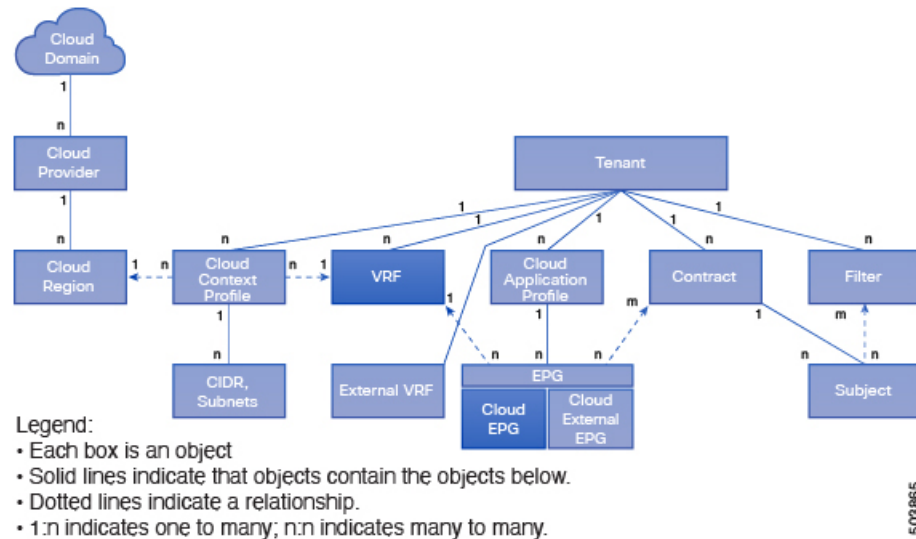
For example, the dotted line between the cloud EPG and the VRF defines the relation between those two MOs. In this figure, the cloud EPG (`cloudEPg`) contains a relationship MO (`cloudRsCloudEPgCtx`) that is named with the name of the target VRF MO (`fvCtx`). For example, if production is the VRF name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

- Google Cloud provider (for the infra tenant)

- Monitoring

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 Cloud Network Controller applies the default policy. An administrator can create a default policy and the Cisco Cloud Network Controller uses that unless the administrator provides any explicit policy.

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.

CHAPTER 5

Configuring Cisco Cloud Network Controller Components

About Configuring the Cisco Cloud Network Controller

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

Note For information about the GUI, such as navigation and a list of configurable components, see About the Cisco Cloud Network Controller GUI, on page 4.

Configuring the Cisco Cloud Network Controller Using the GUI

Creating a Tenant

The following sections describe how to create a managed tenant or unmanaged tenant.

Setting Up the Google Cloud Project for a User Tenant

Perform the procedures in this section to set up the Google Cloud project for a user tenant, where that user tenant is either a managed or an unmanaged tenant.

Step 1 Create a Google Cloud project for the user tenant, if necessary.

Each user tenant is mapped one-to-one to a Google Cloud project. If you do not have a Google Cloud project created yet for your user tenant, follow these procedures to create a Google Cloud project.

a) Log into your Google account.

b) Navigate to **IAM & Admin** > **Manage resources**.

c) Using the **Select organization** drop-down list at the top of the page, choose the organization where you want to create a project.

d) Click + **CREATE PROJECT**.

e) In the **New Project** window that appears, enter a project name and select a billing account as applicable.

   A project name can contain only letters, numbers, single quotes, hyphens, spaces, or exclamation points, and must be between 4 and 30 characters.

f) Enter the parent organization or folder in the **Location** field.

   That resource will be the hierarchical parent of the new project.

g) Click **CREATE**.

Step 2 In Google Cloud, enable the appropriate service APIs in the service account associated with this user tenant.

a) In the Google Cloud GUI, log into the Google Cloud project that is associated with this user tenant.
   The **Dashboard** for the project is displayed.

b) In the search bar at the top of the **Dashboard**, search for **APIs & Services**, then click the result from that search to access the **APIs & Services** window.

c) In the **APIs & Services** window, click the + **ENABLE APIS AND SERVICES** tab.

   The **API Library** window appears.

d) In the **Search for APIs & Services** field, search for and enable the necessary services.

   For each of the services in the list below:

   1. Search for the API or service in the **Search for APIs & Services** field.

   2. Click on the search result to display the page for that API or service.

   3. Click the **ENABLE** button in that API or service page.

   Following are the APIs and services that you must search for and enable:

   - Compute Engine API

   - Cloud Deployment Manager V2 API

   - Cloud Pub/Sub API

   - Cloud Resource Manager API

   - Service Usage API

   - Cloud Logging API

   Each API or service takes several minutes to enable. You will have to navigate back to the **APIs & Services** window after you enable each API or service.

   Note that the following additional APIs and services should be enabled automatically when you enable all of the APIs and services listed above:

   - Identity and Access Management (IAM) API

   - IAM Service Account Credentials API

> - Cloud OS Login API
>
> - Cloud DNS API
>
> - Recommender API
>
> If they are not enabled automatically, enable them manually.

Step 3 Set the necessary permissions for this user tenant in Google Cloud.

a) In the Google Cloud GUI, log into the Google Cloud project that is associated with this user tenant.
The **Dashboard** for the project is displayed.

b) In the left nav bar, click on **IAM & Admin**, then choose **IAM**.

The **IAM** window appears with several service accounts displayed.

c) Locate the appropriate service account.

d) Set the permissions for this service account.

1. Click the pencil icon on the row for this service account.

The **Edit Permissions** window is displayed.

2. Click + **ADD ANOTHER ROLE**, then choose **Editor** as the role.

You are returned to the **IAM** window with the service accounts displayed.

3. Click + **ADD ANOTHER ROLE** again, then add the remaining necessary roles for this service account.

Following is the full list of roles that you must assign to this service account, including the Cloud Functions
Service Agent that you added in the first step of this process:

- Editor

- Role Admin

- Project IAM Admin

4. After you have added all the necessary roles, click **SAVE**.

You are returned to the **IAM** window with the service accounts displayed and the necessary roles assigned to
this service account.

Creating a Managed Tenant

The following sections provide the information that you'll need to create a managed tenant, where you will:

- Create a managed tenant in Cisco Cloud Network Controller

- Set the necessary permissions for the managed tenant in Google Cloud

Creating a Managed Tenant Using the Cisco Cloud Network Controller GUI

This section explains how to create a tenant that will be managed by Cisco Cloud Network Controller using
the GUI.

Step 1 Set up the Google Cloud project for the user tenant.

See Setting Up the Google Cloud Project for a User Tenant, on page 39 for those procedures.

Step 2 In the Cisco Cloud Network Controller GUI, navigate to **Application Management** > **Tenants**.

A table of already-configured tenants is displayed.

Step 3 Click **Actions** and choose **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 3: Create Tenant Dialog Box Fields

Properties	Description
Name	Enter the name of the tenant. Match the regular expression: `[a-z]([-a-z0-9]*[a-z0-9])?` This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen.
Description	Enter a description of the tenant.
Settings	
Add Security Domain	To add a security domain for the tenant: a. Click **Add Security Domain**. The **Select Security Domains** dialog appears with a list of security domains in the left pane. b. Click to choose a security domain. c. Click **Select** to add the security domain to the tenant.
Google Cloud Project	
Google Cloud Project ID	Enter the Google Cloud Project ID that will be associated with this Cisco Cloud Network Controller tenant.
Access Type	For a tenant that will be managed by the Cisco Cloud Network Controller, choose **Managed Identity** as the access type. For more information, see Understanding Google Cloud Deployments with Cisco Cloud Network Controller, on page 6.

Properties	Description
Add Security Domain for Google Cloud Project	**Note** Adding a security domain for Google Cloud is optional when creating a tenant.
	To add a security domain for the account:
	a. Click **Add Security Domain for Google Cloud Project**. The **Select Security Domains** dialog appears with a list of security domains in the left pane.
	b. Click to choose a security domain.
	c. Click **Select** to add the security domain to the tenant.

Step 5   Click **Save** when finished.

What to do next

Complete the necessary configurations in Google Cloud for the managed tenant. Go to for those procedures.

Setting the Necessary Permissions in Google Cloud for a Managed Tenant

If you are creating a managed tenant, you must now set the necessary permissions in Google Cloud.

✎

Note   You do not have to follow the steps in this procedure if you are creating an unmanaged tenant.

Step 1   In the Google Cloud GUI, log into the Google Cloud project that is associated with this managed tenant.

The **Dashboard** for the project is displayed.

Step 2   In the left nav bar, click on **IAM & Admin**, then choose **IAM**.

The **IAM** window appears with several service accounts displayed.

Step 3   Locate the service account that was created in the project that is associated with the infra account.

Step 4   Copy the service account name.

Step 5   Add this service account name as an IAM user in the user tenant project.

Step 6   Set the permissions for this service account.

a)   Click the pencil icon on the row for this service account.

The **Edit Permissions** window is displayed.

b)   Click + **ADD ANOTHER ROLE**, then choose **Cloud Functions Service Agent** as the role.

You are returned to the **IAM** window with the service accounts displayed.

c)   Click + **ADD ANOTHER ROLE** again, then add the remaining necessary roles for this service account.

Following is the full list of roles that you must assign to this service account, including the Cloud Functions Service Agent that you added in the first step of this process:

- Cloud Functions Service Agent

- Compute Instance Admin (v1)

- Compute Network Admin

- Compute Security Admin

- Logging Admin

- Pub/Sub Admin

- Storage Admin

d) After you have added all the necessary roles, click **SAVE**.

You are returned to the **IAM** window with the service accounts displayed and the necessary roles assigned to this service account.

Creating an Unmanaged Tenant

The following sections provide the information that you'll need to create an unmanaged tenant, where you will:

- Generate and download the necessary private key information from Google Cloud for an unmanaged tenant

- Create an unmanaged tenant in Cisco Cloud Network Controller

Generating and Downloading Private Key Information from Google Cloud for an Unmanaged Tenant

If you are creating an unmanaged tenant, you must first generate and download the necessary private key information from Google Cloud.

Note You do not have to follow the steps in this procedure if you are creating a managed tenant.

Step 1 In Google Cloud, select the Google Cloud project that will be associated with this unmanaged tenant, if you have not selected it already .

Step 2 In the left nav bar, click on **IAM & Admin**, then choose **Service Accounts**.

The service accounts for this Google Cloud project are displayed.

Step 3 Select an existing service account or click + **CREATE SERVICE ACCOUNT** to create a new one.

Information on this service account is displayed, with the **Details** tab selected by default.

Step 4 Click the **KEYS** tab.

Step 5 Click **ADD KEY** > **Create New Key**.

A window appears, providing an option to create a private key for this service account.

Step 6   Leave the **JSON** key type selected, then click **Create**.

A window appears, saying that the private key has been saved to your computer.

Step 7   Locate the JSON file that was downloaded to your computer and move it to a secure location on your computer.

This JSON file will contain the key information that you need to fill in the fields for the unmanaged tenant.

```
{
"type": "service_account",
"project_id": "                           ",
"private_key_id": "                                    ",
"private_key": "-----BEGIN PRIVATE
KEY-----




"client_id": "                        ",
"auth_uri": "https://accounts.google.com/o/oauth2/auth",
"token_uri": "https://oauth2.googleapis.com/token",
"auth_provider_x509_cert_url": "https://www.googleapis.com/oauth2/v1/certs",
"client_x509_cert_url": "                                                          "
}
```

Creating an Unmanaged Tenant Using the Cisco Cloud Network Controller GUI

This section explains how to create a tenant that will not be managed by Cisco Cloud Network Controller using the GUI.

Before you begin

Complete the procedures provided in Generating and Downloading Private Key Information from Google Cloud for an Unmanaged Tenant, on page 44 before proceeding with the procedures in this section.

Step 1   Set up the Google Cloud project for the user tenant.

See Setting Up the Google Cloud Project for a User Tenant, on page 39 for those procedures.

Step 2   In the Cisco Cloud Network Controller GUI, navigate to **Application Management** > **Tenants**.

A table of already-configured tenants is displayed.

Step 3   Click **Actions** and choose **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 4: Create Tenant Dialog Box Fields

Properties	Description
Name	Enter the name of the tenant. Match the regular expression: `[a-z]([-a-z0-9]*[a-z0-9])?` This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen.
Description	Enter a description of the tenant.
Settings	
Add Security Domain	To add a security domain for the tenant: a. Click **Add Security Domain**. The **Select Security Domains** dialog appears with a list of security domains in the left pane. b. Click to choose a security domain. c. Click **Select** to add the security domain to the tenant.
Google Cloud Project	
Google Cloud Project ID	Enter the Google Cloud Project ID that will be associated with this Cisco Cloud Network Controller tenant.
Access Type	For a tenant that will not be managed by the Cisco Cloud Network Controller, choose **Unmanaged Identity** as the access type. For more information, see Understanding Google Cloud Deployments with Cisco Cloud Network Controller, on page 6.
Key ID	Enter the information from the `private_key_id` field in the JSON file that you downloaded in Generating and Downloading Private Key Information from Google Cloud for an Unmanaged Tenant, on page 44.
RSA Private Key	Enter the information from the `private_key` field in the JSON file that you downloaded in Generating and Downloading Private Key Information from Google Cloud for an Unmanaged Tenant, on page 44.
Client ID	Enter the information from the `client_id` field in the JSON file that you downloaded in Generating and Downloading Private Key Information from Google Cloud for an Unmanaged Tenant, on page 44.

Properties	Description
Email	Enter the email address associated with your Google Cloud project.
Add Security Domain for Google Cloud Project	**Note** Adding a security domain for Google Cloud is optional when creating a tenant. To add a security domain for the account: **a.** Click **Add Security Domain for Google Cloud Project**. The **Select Security Domains** dialog appears with a list of security domains in the left pane. **b.** Click to choose a security domain. **c.** Click **Select** to add the security domain to the tenant.

Step 5   Click **Save** when finished.

Creating an Application Profile Using the Cisco Cloud Network Controller GUI

This section explains how to create an application profile using the Cisco Cloud Network Controller 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.

Note the following restrictions:

- Match the regular expression:

```
[a-z]([-a-z0-9]*[a-z0-9])?
```

This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen.

- We recommend using 14 characters or fewer for this name, if possible, due to the restrictions imposed by the Google Cloud firewall rules. Refer to Naming Length Restrictions Imposed By Google Cloud Firewall Rules, on page 23 to better understand the restriction and the total number of characters allowed for each of the Cisco Cloud Network Controller components that make up the firewall rule name.

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 Network Controller GUI

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

Note To configure a external VRF, you will select `infra` in the **Tenant** field below. The VRF will be identified as a external VRF if it is:

- Configured under the `infra` tenant

- Associated with an external network (see Creating an External Network Using Cloud Native Routers Using the Cisco Cloud Network Controller GUI, on page 49)

- Not associated with a cloud context profile

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 5: Create VRF Dialog Box Fields

Properties	Description
General	

Properties	Description
Name	Enter a name for the VRF in the **Name** field. Note the following restrictions: • Match the regular expression: `[a-z]([-a-z0-9]*[a-z0-9])?` This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen. • We recommend using 14 characters or fewer for this name, if possible, due to the restrictions imposed by the Google Cloud firewall rules. Refer to Naming Length Restrictions Imposed By Google Cloud Firewall Rules, on page 23 to better understand the restriction and the total number of characters allowed for each of the Cisco Cloud Network Controller components that make up the firewall rule name. 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. To see the *vrfEncoded* value, navigate to **Application Management** > **VRFs** subtab. Click a VRF on the right hand pane and look for *Encoded VRF Name in Cloud Router*.
Tenant	To choose a tenant: a. Click **Select Tenant**. The **Select Tenant** dialog box appears. b. From the **Select Tenant** dialog, click to choose a tenant in the left column then click **Select**. You return to the **Create VRF** dialog box.
Description	Enter a description of the VRF.

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

Creating an External Network Using Cloud Native Routers Using the Cisco Cloud Network Controller GUI

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

Before you begin

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

Step 1   In the left navigation bar, navigate to **Application Management** > **External Networks**.

The configured external networks are displayed. Note that because Cisco Cloud Network Controller supports only one hub network, you will see only one hub network displayed in the **Hub Network** column.

Step 2 Click **Actions**, then choose **Create External Network**.

The **Create External Network** window appears.

Note If there is no hub network configured yet, you will see a warning at the top of the page, saying that you must create a hub network before you can create an external network. Click the blue `Cisco Cloud Network Controller Setup` link in the message to create a hub network, then return here. For more information on creating a hub network, see the "Configuring Cisco Cloud Network Controller Using the Setup Wizard" chapter in the Cisco Cloud Network Controller for Google Cloud Installation Guide, Release 25.0(x) or later.

Step 3 Enter the appropriate values in each field as listed in the following *Create External Network Dialog Box Fields* table then continue.

Table 6: Create External Network Dialog Box Fields

Properties	Description
General	
Name	Enter the name for the external network.
VRF	This external VRF will be used for external connectivity with the on-premises CCR. You can create multiple external VRFs for this purpose.
	This VRF will be identified as an external VRF if the VRF has all three of the following characteristics:
	• Configured under the `infra` tenant
	• Associated with an external network
	• Not associated with a cloud context profile
	Any VRF that is associated with an external network becomes an external VRF. At that point, that external VRF is not allowed to be created under any tenant other than the infra tenant, and that external VRF is not allowed to be associated with a cloud context profile or subnet.
	To choose an external VRF:
	a. Click **Select VRF**.
	The **Select VRF** dialog box appears.
	b. From the **Select VRF** dialog, click to choose a VRF in the left column.
	You can also create a VRF using the + **Create VRF** option.
	c. Click **Select**.
	You return to the **Create External Network** dialog box.
Hub Network	The hub network is displayed automatically after you configured it in the First Time Setup.
	Note If there is no hub network configured yet, you must create a hub network before you can create an external network. For more information on creating a hub network, see the "Configuring Cisco Cloud Network Controller Using the Setup Wizard" chapter in the Cisco Cloud Network Controller for Google Cloud Installation Guide, Release 25.0(x) or later.

Properties	Description
VPN Router	This field is not editable. The default VPN router is automatically selected.
Settings	

Regions	To choose a region:
	
	a. Click **Add Regions**.
	
	    The **Select Regions** dialog box appears.
	
	      • The regions that you selected as part of the First Time Setup are displayed here.
	
	      • You can select multiple regions to bring up the cloud router in multiple regions.
	
	b. From the **Select Regions** dialog, click to choose a region in the left column then click **Select**.
	
	    You return to the **Create External Network** dialog box.

Properties	Description

VPN Networks	The VPN networks entries are used for internal connectivity. All configured VPN networks will be applied to all the selected regions.
	
	To add a VPN network:
	
	a. Click **Add VPN Network**.
	
	The **Add VPN Network** dialog box appears.
	
	b. In the **Name** field, enter a name for the VPN network.
	
	c. Click + **Add IPSec Peer**.
	
	Two tunnels are created for each IPSec peer entry.
	
	d. Enter values for the following fields for the IPSec peer that you want to add:
	
	    • **Public IP of IPSec Tunnel Peer**
	
	    • **Pre-Shared Key**
	
	    • **IKE Version**: Select **ikev1** or **ikev2** for IPSec tunnel connectivity
	
	    • **BGP Peer ASN**
	
	    • **Subnet Pool Name**: Click **Select Subnet Pool Name**.
	
	    The **Select Subnet Pool Name** dialog box appears. Select one of the available subnet pools that are listed, then click **Select**.
	
	    Note Additional IPsec tunnel subnet pools can be added in the **External Networks** page, or through the Cloud Network Controller First Time Set Up, if necessary. For more information on adding additional subnet pools through the Cloud Network Controller First Time Set Up, see the chapter "Configuring Cisco Cloud Network Controller Using the Setup Wizard" in the *Cisco Cloud Network Controller for GCP Installation Guide*, Release 25.0(5). The subnet pool size should be large enough to accommodate the number of IPsec tunnels that will get created.
	
	e. Click the checkmark to add this IPSec tunnel.
	
	Click + **Add IPSec Tunnel** if you want to add another IPSec tunnel.
	
	f. Click **Add** in the **Add VPN Network** dialog box.
	
	You return to the **Create External Network** dialog box.

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

To verify that cloud routers were configured in Google Cloud, in your Google Cloud account, navigate to **Hybrid Connectivity** > **Cloud Routers**. You should see the cloud routers created for the different regions (note that you might have to click Refresh to bring up the newly-configured cloud routers).

To see the IPSec sessions, navigate to **Hybrid Connectivity** > **VPN** > **Cloud VPN Tunnels**.

Configuring Inter-Site Networking Using BGP-EVPN

Beginning with release 25.0(5), for inter-site use cases, support is available for configuring a BGP-EVPN connection for inter-site connectivity in these scenarios:

- Cloud site-to-cloud site:

    - Google Cloud site-to-Google Cloud site

    - Google Cloud site-to-AWS site

    - Google Cloud site-to-Azure site

- Google Cloud site-to-ACI on-premises site

In each of these scenarios, Cisco Catalyst 8000Vs are used for the BGP-EVPN connection. See Inter-Site Connectivity Using BGP-EVPN, on page 12 for more information.

Step 1 In the Cisco Cloud Network Controller GUI, click the Intent icon (🔵) and select **Cloud Network Controller Setup**.

Step 2 In the **Region Management** area, click **Edit Configuration**.

The **Region Management** page appears.

Step 3 Locate the regions where you want to use the Cisco Catalyst 8000V router for inter-site connectivity, then click the box in the **Catalyst 8000Vs** column for those regions.

This is functionality introduced in release 25.0(5) that allows you to configure a BGP-EVPN connection for inter-site connectivity between a Google Cloud site and other cloud sites or an ACI on-premises site using Cisco Catalyst 8000V routers. See Inter-Site Connectivity Using BGP-EVPN, on page 12 for more information.

Step 4 Click **Next** at the bottom of the page.

The **General Connectivity** page appears.

Step 5 Enter the necessary information in the **Subnet Pools for Cloud Routers** area.

The first subnet pool is automatically populated (shown as System Internal). Addresses from this subnet pool will be used for inter-region connectivity for any additional regions that are added that need to be managed by the Cisco Cloud Network Controller. Subnet pools added in this field must be a valid IPv4 subnet with mask /24.

If you selected additional regions to have Catalyst 8000Vs deployed in the previous page, add *one* additional subnet pool for every region where you will have 2-4 Catalyst 8000Vs deployed (if you enter **2**, **3**, or **4** in the **Number of Routers Per Region** field in 6.c, on page 54).

Step 6 For release 25.0(5) and later, enter the necessary information in the **Catalyst 8000Vs** area.

a) In the **BGP Autonomous System Number for C8kVs** field, enter a unique BGP autonomous system number (ASN).

The BGP autonomous system number can be in the range of 1 - 65535.

b) In the **Assign Public IP to C8kV Interface** field, determine if you want to assign public IP addresses to the Catalyst 8000V interfaces.

Private IP addresses are assigned to the Catalyst 8000V interfaces by default. The **Assign Public IP to C8kV Interface** option determines whether public IP addresses will also be assigned to the Catalyst 8000V interfaces or not.

The Catalyst 8000V interface IP addresses are used for the following purposes:

- Allows you to manage the Catalyst 8000V or allows you to SSH to the Catalyst 8000V directly

- Allows you to cross-program the interfaces across sites for multi-cloud and hybrid cloud connectivity through the Cisco Nexus Dashboard Orchestrator

- For the Catalyst 8000Vs for both control plane and data plane traffic

By default, the **Enabled** check box is checked. This means that public IP addresses can be assigned to the Catalyst 8000Vs.

- If you want *public* IP addresses assigned to the Catalyst 8000Vs in addition to the private IP addresses, leave the check in the box next to **Enabled**.

- If you want only *private* IP addresses assigned to the Catalyst 8000Vs, remove the check in the box next to **Enabled** to disable this option.

Note that changing the Catalyst 8000V connectivity from private to public, or vice versa, may cause disruption in your network. In addition, if the public IP address is removed from the Catalyst 8000V, then the Google Cloud site will connect to the on-premises ACI site using the private IP address via the Google Cloud interconnect. You will have to configure private intersite connectivity for the Google Cloud site from Nexus Dashboard Orchestrator and configure Google Cloud interconnect from the Google Cloud portal.

Note Both the public and private IP addresses assigned to a Catalyst 8000V are displayed with the other details of the router in the **Cloud Resources** area. If public IP addresses are not assigned to a Catalyst 8000V, only the private IP addresses are displayed.

c) In the **Number of Routers Per Region** field, choose the number of Catalyst 8000Vs that will be used in each region.

d) In the **Username**, enter the username for the Catalyst 8000V.

e) In the **Password** field, enter the password for the Catalyst 8000V.

Enter the password again in the **Confirm Password** field.

f) In the **Throughput of the routers** field, choose the throughput of the Catalyst 8000V.

Changing the value in this field changes the size of the Catalyst 8000V instance that is deployed. Choosing a higher value for the throughput results in a larger VM being deployed.

Note the following:

- The licensing of the Catalyst 8000V is based on this setting. You will need the equivalent or higher license in your Smart account for it to be compliant. See "Resources Used for Cisco Cloud Network Controller Deployment in Google Cloud" in the *Cisco Cloud Network Controller for Google Cloud Installation Guide* for more information.

- Cloud routers should be undeployed from all regions before changing the router throughput or login credentials.

If you wish to change this value at some point in the future, you must delete the Catalyst 8000V, then repeat the processes in this chapter again and select the new value that you would like in the same **Throughput of the routers** field.

g) Enter the necessary information in the **TCP MSS** field, if applicable.

The **TCP MSS** option is available to configure the TCP maximum segment size (MSS). This value will be applied to all cloud router interfaces, including data Gigabit Ethernet interfaces, IPSec tunnel interfaces of cloud routers,

and VPN tunnel interfaces toward cloud, on-premises, or other cloud sites. For VPN tunnels towards the cloud, if the cloud provider's MSS value is less than the value that you enter in this field, then the lower value is used; otherwise, the value that you enter in this field is used.

The MSS value affects only TCP traffic, and has no impact on other types of traffic, such as ping traffic.

h) In the **License Token** field, enter the license token for the Catalyst 8000V.

This is the Product Instance Registration token from your Cisco Smart Software Licensing account. To get this license token, go to http://software.cisco.com, then navigate to **Smart Software Licensing** > **Inventory** > **Virtual Account** to find the Product Instance Registration token. See "Cisco Cloud Network Controller Licensing" in the *Cisco Cloud Network Controller for Google Cloud Installation Guide* for more information.

Note   If you assigned private IP addresses to the Catalyst 8000Vs in 6.b, on page 53, the only supported option is **Direct connect to Cisco Smart Software Manager (CSSM)** when registering smart licensing for Catalyst 8000Vs with private IP addresses. You must provide reachability to the CSSM through express route in this case.

Step 7   When you have entered all the necessary information on this page, click **Save and Continue** at the bottom of the page.

Step 8   In the **Advanced Settings** area, click **Edit Configuration**.

Step 9   In the **Contract Based Routing** field, click the box next to **yes** to enable contract-based routing, then click **Save and Continue**.

Note   You can also enable contract-based routing through Nexus Dashboard Orchestrator by navigating to the Google Cloud site in Nexus Dashboard Orchestrator, then clicking the **Contract Based Routing** option under the **Inter-Site Connectivity** area.

Step 10   Click **Done** at the bottom of the **Let's Configure the Basics** window.

Step 11   Verify that the number of VM instances on the Google Cloud side match the number of Catalyst 8000Vs that you set up in the Cisco Cloud Network Controller.

a) Log into the Google Cloud project associated with the infra tenant.

b) Navigate to **Compute Engine** > **VM instances** in Google Cloud.

c) Verify that the number of VM instances shown in the **Instances** tab match the total number of Catalyst 8000Vs that you have for the BGP-EVPN connection for inter-site connectivity.

For example, if you chose two regions and two Catalyst 8000Vs for each region, you should see four VM instances in the **Instances** tab.

Step 12   Verify that you have the VPC networks set up for the overlay-1 VPC and overlay-1 secondary VPC in Google Cloud.

See Inter-Site Connectivity Using BGP-EVPN, on page 12 for more information.

a) Navigate to **VPC network** > **VPC networks** in Google Cloud.

b) Verify that you see the VPC networks that were set up for the overlay-1 VPC and overlay-1 secondary VPC in the **VPC networks** screen.

Step 13   Verify that the Catalyst 8000Vs that you set up in the Cisco Cloud Network Controller are showing up correctly in the appropriate Cisco Cloud Network Controller GUI screens.

• In the **Dashboard** page, locate the **Connectivity** pane and verify that the Catalyst 8000Vs that you set up in the Cisco Cloud Network Controller show up correctly in this screen.

• Navigate to **Infrastructure** > **Inter-Site Connectivity** and verify that the Catalyst 8000Vs that you set up in the Cisco Cloud Network Controller show up correctly in this screen.

> • Navigate to **Cloud Resources** > **Routers** and verify that the Catalyst 8000Vs that you set up in the Cisco Cloud Network Controller show up correctly in this screen.

Step 14 Configure VPC peering for the inter-site connectivity using BGP-EVPN.

When you configure a BGP-EVPN connection for inter-site connectivity using Cisco Catalyst 8000V routers, you must make additional configurations to allow the user VPCs in the Google Cloud site to communicate with VPCs in other cloud sites or an ACI on-premises site.

Typically you would configure VPC peering for inter-site connectivity using BGP-EVPN through Nexus Dashboard Orchestrator, where you would create a VRF and then check **Hub Peering** for that VRF. See the appropriate Nexus Dashboard Orchestrator documentation for those procedures.

To change this configuration on the Cisco Cloud Network Controller side:

a) In the Cisco Cloud Network Controller GUI, navigate to **Application Management** > **Cloud Context Profiles**.
b) Under the **Name** column, double-click the name of the cloud context profile that is associated with the VPC that you want to peer with the overlay-1 VPC.

   Another window appears that provides more detailed information for this cloud context profile.

c) Click **Actions** > **Edit**.
d) In the **VPC Hub Peering** area, click the box next to **Enable** to enable VPC peering for this VPC, then click **Save**.
e) In the Google Cloud, navigate to **VPC network** > **VPC network peering**.
f) Verify that your user VPC in the Google Cloud site is peering with the overlay-1 VPC.

Configuring Inter-VRF Route Leaking Using the Cisco Cloud Network Controller GUI

Using inter-VRF route leaking, you can configure an independent routing policy to specify which routes to leak between a pair of VRFs when you are setting up routing between these types of sites:

• Two cloud sites

• A cloud site and a non-ACI on-premises site

> **Note** See Configuring Routing and Security Policies Separately, on page 14 for more information.

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

Table 7: Create Leak Routes Dialog Box Fields

Properties	Description
Source VRF	To choose a source VRF: a. Click **Select a Source VRF**. The **Select a VRF** dialog box appears. b. From the **Select a VRF** dialog, click to choose a VRF in the left column to use for the source VRF. Note that the source VRF can be an internal or an external (transport) VRF. c. Click **Select** to select this source VRF. You return to the **Create Leak Route** dialog box.
Destination VRF	To choose a destination VRF: a. Click **Select a Destination VRF**. The **Select a VRF** dialog box appears. b. From the **Select a VRF** dialog, click to choose a VRF in the left column to use for the destination VRF. c. Click **Select** to select this destination VRF. You return to the **Create Leak Route** dialog box.
Type	Choose the type of leaked route that you want to configure: • **Leak All**: Select to configure all routes to leak between the VRFs. The entry `0.0.0.0/0` is entered automatically in the subnet IP area by default in this case. • **Subnet IP**: Select to configure a specific subnet IP address as the route to leak between VRFs. The **Subnet IP** box appears. In the **Subnet IP** box, enter a subnet IP address as the route to leak between VRFs. To configure multiple subnet IP addresses as the route to leak between VRFs, enter additional entries for the different subnets.

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

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

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

  You are returned to the **Add Leak Route** window. Repeat through to configure another route to leak between a pair of VRFs.

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

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

     • The source VRF from the previous configuration now becomes the destination VRF

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

You are returned to the **Add Leak Route** window. Repeat Step 4, on page 56 through Step 5, on page 57 to configure another route, but this time:

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

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

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

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

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

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

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

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

     The **Add Leak Route** window appears. Enter the necessary information as you did previously using the information in Step 4, on page 56. Note that the entry in the **Source VRF** is pre-selected and cannot be changed in this situation.

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

     The **Add Leak Route** window appears. Enter the necessary information as you did previously using the information in Step 4, on page 56. Note that the entry in the **Destination VRF** is pre-selected and cannot be changed in this situation.

Enabling Connectivity Between Google Cloud and External Devices

Follow these procedures to manually enable connectivity between a Google Cloud Router and an external device.

Downloading the External Device Configuration Files

Step 1   In the Cisco Cloud Network Controller GUI, click on **Dashboard**.
The **Dashboard** view for the Cisco Cloud Network Controller appears.

Step 2   In the **Connectivity** area, under **External Connectivity Status**, click on the number above the **Cloud Routers** entry.
The **External Connectivity** window appears.

Step 3   Click **Actions** > **Download External Device Configuration Files**.
The **Download External Device Configuration Files** pop-up appears.

Step 4   Select the external device configuration files to download and click **Download**.

This action downloads a zip file that contains configuration information that you will use to enable connectivity between the Google Cloud Router and the external devices.

Enabling Connectivity Between Google Cloud and the External Devices

Before you begin

Download the external device configuration files using the procedures in Downloading the External Device Configuration Files, on page 58.

Step 1 Gather the necessary information that you will need to enable connectivity between the Google Cloud Router and the external devices.

Step 2 Log into the external device.

Step 3 Enter the configuration information to connect an external networking device with the cloud ACI fabric.

If you downloaded the external device configuration files using the instructions in Downloading the External Device Configuration Files, on page 58, locate the configuration information for the first tunnel and enter that configuration information.

Following is an example of what the external device configuration file might look like for the first tunnel, where **PRESHARED-KEY** is taken from the vpn-connectivity configuration page:

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

! Tunnel to 54.215.245.58 5.500 for
hcextnwTunnIf.acct-[infra]/region-[us-west1]/hubCtx-[1]-id-[0]/ext-[extnwfoo_us-west1]/vpn-[vpnnwfoo]/rtr-default-peer-54.215.245.58/src-1-dest-[54.215.245.58]
! USER-DEFINED: please define rd: RD
! USER-DEFINED: please provide preshared-key: PRESHARED-KEY
! USER-DEFINED: please define router-id: ROUTER-ID
! USER-DEFINED: please define gig-number: GIG-NUMBER
! USER-DEFINED: please define gig-gateway: GIG-GATEWAY
! ikev: ikev2
! vrf-name: extv1
! user name: root
! tunnel counter: 5
! IPV4 address: 35.220.50.132
! tunnel interface destination: 54.215.245.58
! tunne id: 500
! BGP peer address: 169.254.10.6
! BGP peer neighbor address: 169.254.10.5
! BGP peer ASN: 64513
! hcloudHubCtx ASN: 64512

vrf definition extv1
    rd RD:1
    address-family ipv4
    exit-address-family
exit

interface Loopback0
    vrf forwarding extv1
    ip address 41.41.41.41 255.255.255.255
```

```
    exit


crypto ikev2 proposal ikev2-1
    encryption aes-cbc-256 aes-cbc-192 aes-cbc-128
    integrity sha512 sha384 sha256 sha1
    group 24 21 20 19 16 15 14 2
exit

crypto ikev2 policy ikev2-1
    proposal ikev2-1
exit

crypto ikev2 keyring keyring-root-5
    peer peer-ikev2-keyring
        address 35.220.50.132
        pre-shared-key PRESHARED-KEY
    exit
exit

crypto ikev2 profile ikev-profile-root-5
    match address local interface GIG-NUMBER
    match identity remote address 35.220.50.132 255.255.255.255
    identity local address 54.215.245.58
    authentication remote pre-share
    authentication local pre-share
    keyring local keyring-root-5
    lifetime 3600
    dpd 10 5 periodic
exit

crypto ipsec transform-set ikev-transport-root-5 esp-gcm 256
    mode tunnel
exit

crypto ipsec profile ikev-profile-root-5
    set transform-set ikev-transport-root-5
    set pfs group14
    set ikev2-profile ikev-profile-root-5
exit

interface Tunnel500
    vrf forwarding extv1
    ip address 169.254.10.6 255.255.255.252
    ip mtu 1400
    ip tcp adjust-mss 1400
    tunnel source GIG-NUMBER
    tunnel mode ipsec ipv4
    tunnel destination 35.220.50.132
    tunnel protection ipsec profile ikev-profile-root-5
exit

ip route 35.220.50.132 255.255.255.255 GIG-NUMBER GIG-GATEWAY

router bgp 64513
    bgp router-id ROUTER-ID
    bgp log-neighbor-changes

    address-family ipv4 vrf extv1
        network 41.41.41.41 mask 255.255.255.255
        neighbor 169.254.10.5 remote-as 64512
        neighbor 169.254.10.5 ebgp-multihop 255
        neighbor 169.254.10.5 activate
    exit-address-family
```

```
exit
```

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

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

    - VRF definition

    - IPSec global configurations

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

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

    - IPSec and ikev1 per tunnel configurations

    - BGP configurations for the VRF neighbor

```
!
crypto keyring Ext-V1-1000-ike
  pre-shared-key address <50.18.55.126>[cAPIC CSR Gig3 public IP] key <abcdefg12345>
!
crypto isakmp profile Ext-V1-1000-ike
    keyring Ext-V1-1000-ike
    match identity address <50.18.55.126>[cAPIC CSR1 gig3 Public IP] 255.255.255.255
!
crypto ipsec transform-set Ext-V1-1000-ike esp-aes esp-sha-hmac
mode tunnel
!
crypto ipsec profile Ext-V1-1000-ike
set security-association lifetime kilobytes disable
set security-association replay window-size 512
set transform-set Ext-V1-1000-ike
set pfs group14
!
interface Tunnel1000
vrf forwarding Ext-V1
ip address 50.50.0.2[cAPIC CSR BGP Peer Addr] 255.255.255.252
ip mtu 1400
ip tcp adjust-mss 1400
tunnel source GigabitEthernet2
tunnel mode ipsec ipv4
tunnel destination <50.18.55.126>[cAPIC CSR1 gig3 Public IP]
tunnel protection ipsec profile Ext-V1-1000-ike
!
router bgp 64550
!
address-family ipv4 vrf Ext-V1
  redistribute connected
  neighbor <50.50.0.1>[cAPIC CSR1 Tunnel Inner IP Addr] remote-as 1234
  neighbor 50.50.0.1 ebgp-multihop 255
  neighbor 50.50.0.1 activate
  neighbor 50.50.0.1 send-community both
  neighbor <50.50.0.5>[cAPIC CSR1 Tunnel Inner IP Addr] remote-as 1234
  neighbor 50.50.0.5 ebgp-multihop 255
  neighbor 50.50.0.5 activate
  neighbor 50.50.0.5 send-community both
  distance bgp 20 200 20
!
ip route 50.18.55.126[cAPIC CSR1 gig3 Public IP] 255.255.255.255 GigabitEthernet2 10.10.0.103
```

IPSec and Ikev1
Per Tunnel Configurations

BGP Configurations for VRF Neighbor

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

    - Ikev2 global configurations

    - IPSec and ikev2 per tunnel configurations

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

Ikev2 Global Configurations

IPSec and Ikev2
Per Tunnel Configurations

Creating an EPG Using the Cisco Cloud Network Controller GUI

Use the procedures in this section to create an application EPG or an external EPG. The available configuration options vary, depending on which type of EPG you are creating.

Creating an Application EPG Using the Cisco Cloud Network Controller GUI

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

Before you begin

Create an application profile and a VRF.

Step 1   Click the **Intent** icon.

The **Intent** menu appears.

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

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

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

The **Create EPG** dialog box appears.

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

Table 8: Create EPG Dialog Box Fields

Properties	Description
General	
Name	Enter the name of the EPG. Note the following restrictions: • Match the regular expression: `[a-z]([-a-z0-9]*[a-z0-9])?` This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen. • We recommend using 14 characters or fewer for this name, if possible, due to the restrictions imposed by the Google Cloud firewall rules. Refer to Naming Length Restrictions Imposed By Google Cloud Firewall Rules, on page 23 to better understand the restriction and the total number of characters allowed for each of the Cisco Cloud Network Controller components that make up the firewall rule name.
Tenant	To choose a tenant: a. Click **Select Tenant**. The **Select Tenant** dialog box appears. b. From the **Select Tenant** dialog, click to choose a tenant in the left column. c. Click **Select**. You return to the **Create EPG** dialog box.

Properties	Description
Application Profile	To choose an application profile:
	a. Click **Select Application Profile**. The **Select Application Profile** dialog box appears.
	b. From the **Select Application Profile** dialog, click to choose an application profile in the left column.
	Note If you are creating an EPG in the infra tenant, we recommend that you do not choose the `cloud-infra` application profile because that application profile is used by EPGs in the overlay-1 VRF. Select a different application profile or click **Create Application Profile** to create a new one.
	c. Click **Select**. You return to the **Create EPG** dialog box.
Description	Enter a description of the EPG.
Settings	
Type	Because this will be an application EPG, choose **Application** as the EPG type.
VRF	To choose a VRF:
	a. Click **Select VRF**. The **Select VRF** dialog box appears.
	b. From the **Select VRF** dialog, click to choose a VRF in the left column.
	c. Click **Select**. You return to the **Create EPG** dialog box.

Properties	Description
Endpoint Selectors	

Properties	Description
	Note See Configuring Virtual Machines in Google Cloud, on page 79 for instructions on configuring virtual machines in Google Cloud as part of the endpoint selector configuration process.
	To add an endpoint selector:
	a. Click **Add Endpoint Selector** to open the **Add Endpoint Selector** dialog.
	b. In the **Add Endpoint Selector** dialog, enter a name in the **Name** field.
	c. Click **Selector Expression**. The **Key**, **Operator**, and **Value** fields are enabled.
	d. Click the **Key** drop-down list to choose a key. The options are:
	• Choose **IP** if you want to use an IP address or subnet for the endpoint selector.
	• Choose **Region** if you want to use the Google Cloud region for the endpoint selector.
	• Choose **Custom** if you want to create a custom key for the endpoint selector.
	Note When choosing the **Custom** option, the drop-down list becomes a text box. You need to enter a name for the key in the spaces after **custom:** (for example, **custom: Location**).
	e. Click the **Operator** drop-down list to choose an operator. The options are:
	• **equals**: Used when you have a single value in the Value field.
	• **not equals**: Used when you have a single value in the Value field.
	• **in**: Used when you have multiple comma-separated values in the Value field.
	• **not in**: Used when you have multiple comma-separated values in the Value field.
	• **has key**: Used if the expression contains only a key.
	• **does not have key**: Used for an expression that does not contain a key.
	f. Enter a value in the **Value** field then click the check mark to validate the entries. The value you enter depends on the choices you made for the **Key** and **Operator** fields. For example, if the **Key** field is set to **IP** and the **Operator** field is set to **equals**, the **Value** field must be an IP address or subnet. However, if the **Operator** field is set to **has key**, the **Value** field is disabled.
	g. When finished, click the check mark to validate the selector expression.
	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.
	For example, assume you created two sets of expressions under a single endpoint selector:
	• Endpoint selector 1, expression 1:
	• **Key:** Region
	• **Operator:** equals
	• **Value:** us-west1

Properties	Description

	• Endpoint selector 1, expression 2:
	
	    • **Key:** IP
	
	    • **Operator:** equals
	
	    • **Value:** 192.0.2.1/24
	
	In this case, if *both* of these expressions are true (if the region is us-west1 AND if the IP address belongs to subnet 192.0.2.1/24), then that endpoint is assigned to the Cloud EPG.
	
	i. Click the check mark after every additional expression that you want to create under this endpoint selector then click **Add** when finished.
	
	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:
	
	• Endpoint selector 2, expression 1:
	
	    • **Key:** Region
	
	    • **Operator:** in
	
	    • **Value:** us-east1, us-central1
	
	In this case:
	
	• If the region is us-west1 AND the IP address belongs to the 192.0.2.1/24 subnet (endpoint selector 1 expressions)
	
	    OR
	
	• If the region is either us-east1 or us-central1 (endpoint selector 2 expression)
	
	Then that end point is assigned to the Cloud EPG.

Step 5   Click **Save** when finished.

Creating an External EPG Using the Cisco Cloud Network Controller GUI

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

Before you begin

Create an application profile and a VRF.

Step 1   Click the **Intent** icon.

The **Intent** menu appears.

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

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

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

The **Create EPG** dialog box appears.

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

Table 9: Create EPG Dialog Box Fields

Properties	Description
General	
Name	Enter the name of the EPG. Note the following restrictions: • Match the regular expression: `[a-z]([-a-z0-9]*[a-z0-9])?` This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen. • We recommend using 14 characters or fewer for this name, if possible, due to the restrictions imposed by the Google Cloud firewall rules. Refer to Naming Length Restrictions Imposed By Google Cloud Firewall Rules, on page 23 to better understand the restriction and the total number of characters allowed for each of the Cisco Cloud Network Controller components that make up the firewall rule name.
Tenant	To choose a tenant: a. Click **Select Tenant**. The **Select Tenant** dialog box appears. b. From the **Select Tenant** dialog, click to choose a tenant in the left column. c. Click **Select**. You return to the **Create EPG** dialog box.
Application Profile	To choose an application profile: a. Click **Select Application Profile**. The **Select Application Profile** dialog box appears. b. From the **Select Application Profile** dialog, click to choose an application profile in the left column. **Note** If you are creating an EPG in the infra tenant, we recommend that you do not choose the `cloud-infra` application profile because that application profile is used by EPGs in the overlay-1 VRF. Select a different application profile or click **Create Application Profile** to create a new one. c. Click **Select**. You return to the **Create EPG** dialog box.
Description	Enter a description of the EPG.
Settings	

Properties	Description
Type	Because this will be an external EPG, choose **External** as the EPG type.
VRF	To choose a VRF: **a.** Click **Select VRF**. The **Select VRF** dialog box appears. **b.** From the **Select VRF** dialog, click to choose a VRF in the left column. **c.** Click **Select**. You return to the **Create EPG** dialog box.
Route Reachability	The type of route reachability for the external EPG will be automatically selected (either **Internet** or **External-Site**).

Endpoint Selectors	**Note** See Configuring Virtual Machines in Google Cloud, on page 79 for instructions on configuring virtual machines in Google Cloud as part of the endpoint selector configuration process.
	
	To add an endpoint selector:
	
	a. Click **Add Endpoint Selector** to add an endpoint selector.
	
	b. Enter a name in the **Name** field.
	
	c. Enter a subnet in the **Subnet**.
	
	d. When finished, click the check mark to validate the endpoint selector.
	
	e. Determine if you want to create additional endpoint selectors.
	
	If you create more than one endpoint selector under an EPG, a logical OR exists between those endpoint selectors. For example, assume you created two endpoint selectors:
	
	    • Endpoint selector 1:
	
	        • **Name:** EP_Sel_1
	
	        • **Subnet:** 192.1.1.1/24
	
	    • Endpoint selector 2:
	
	        • **Name:** EP_Sel_2
	
	        • **Subnet:** 192.2.2.2/24
	
	In this case:
	
	    • If the IP address belongs to the 192.1.1.1/24 subnet (endpoint selector 1)
	
	      OR
	
	    • If the IP address belongs to the 192.2.2.2/24 subnet (endpoint selector 2)
	
	Then that end point is assigned to the Cloud EPG.

Step 5   Click **Save** when finished.

Creating a Filter Using the Cisco Cloud Network Controller GUI

This section explains how to create a filter using the Cisco Cloud Network Controller 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 10: Create Filter Dialog Box Fields

Properties	Description
Name	Enter a name for the filter in the **Name** field.
Tenant	To choose a tenant: a. Click **Select Tenant**. The **Select Tenant** dialog box appears. b. From the **Select Tenant** dialog, click to choose a tenant in the left column then click **Select**. You return to the **Create Filter** dialog box.
Description	Enter a description of the filter.

Properties	Description
Add Filter	To add a filter:
	a. Click **Add Filter Entry**. The **Add Filter Entry** dialog box appears.
	b. Enter a name for the filter entry in the **Name** field.
	c. Click the **Ethernet Type** drop-down list to choose an ethernet type. The options are:
	• **IP**
	• **Unspecified**
	Note When **Unspecified** is chosen, any traffic type is allowed, including IP, and the remaining fields are disabled.
	d. Click the **IP Protocol** drop-down menu to choose a protocol. The options are:
	• **ICMP**
	• **TCP**
	• **UDP**
	• **Unspecified**
	Note The remaining fields are enabled only when **TCP** or **UDP** is chosen.
	e. Enter the appropriate port range information in the **Destination Port** fields.
	f. When finished entering filter entry information, click **Add**. You return to the **Create Filter** dialog box where you can repeat the steps to add another filter entry.

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

Creating a Contract Using the Cisco Cloud Network Controller GUI

This section explains how to create a contract using the Cisco Cloud Network Controller 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 11: Create Contract Dialog Box Fields

Properties	Description

Name	Enter the name of the contract.
	
	Note the following restrictions:
	
	  • Match the regular expression:
	
	    `[a-z]([-a-z0-9]*[a-z0-9])?`
	
	    This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen.
	
	  • We recommend using 14 characters or fewer for this name, if possible, due to the restrictions imposed by the Google Cloud firewall rules. Refer to Naming Length Restrictions Imposed By Google Cloud Firewall Rules, on page 23 to better understand the restriction and the total number of characters allowed for each of the Cisco Cloud Network Controller components that make up the firewall rule name.

Tenant	To choose a tenant: a. Click **Select Tenant**. The **Select Tenant** dialog box appears. b. From the **Select Tenant** dialog, click to choose a tenant in the left column. c. Click **Select**. You return to the **Create Contract** dialog box.
Description	Enter a description of the contract.
Settings	

Properties	Description
Scope	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.
	Note Shared services enables communication between EPGs in different tenants and between EPGs in different VRFs.
	To enable EPGs in one tenant to communicate with EPGs in another tenant, choose **Global** scope.
	To enable an EPG in one VRF to communicate with another EPG in a different VRF, choose **Global** or **Tenant** scope.
	Click the drop-down arrow to choose from the following scope options:
	• **Application Profile**
	• **VRF**
	• **Global**
	• **Tenant**
Add Filter	To choose a filter:
	a. Click **Add Filter**. The filter row appears with a **Select Filter** option.
	b. Click **Select Filter**. The **Select Filter** dialog box appears.
	c. From the **Select Filter** dialog, click to choose a filter in the left column then click **Select**. You return to the **Create Contract** dialog box.

Step 5   Click **Save** when finished.

Creating an Inter-Tenant Contract Using the Cisco Cloud Network Controller GUI

This section explains how to create an inter-tenant contract using the Cisco Cloud Network Controller 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 12: Create Contract Dialog Box Fields

Properties	Description
Name	Enter the name of the contract.
	This is the name of the contract in Google Cloud. Match the regular expression:
	`[a-z]([-a-z0-9]*[a-z0-9])?`
	This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen.
Tenant	To choose a tenant:
	a. Click **Select Tenant**. The **Select Tenant** dialog box appears.
	b. From the **Select Tenant** dialog, click to choose a tenant in the left column.
	c. Click **Select**. You return to the **Create Contract** dialog box.
Description	Enter a description of the contract.
Settings	
Scope	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.
	For inter-tenant communication, you will first create a contract with the **Global** scope in one of the tenants (for example, **tenant1**). This tenant's EPG will always be the provider of this contract.
	This contract will then be exported to the other tenant (for example, **tenant2**). For the other tenant that imports this contract, its EPG will be the consumer of the imported contract. If you want tenant2's EPG to be the provider and tenant1's EPG to be the consumer, then create a contract in **tenant2** and then export it to **tenant1**.
Add Filter	To choose a filter:
	a. Click **Add Filter**. The filter row appears with a **Select Filter** option.
	b. Click **Select Filter**. The **Select Filter** dialog box appears.
	c. From the **Select Filter** dialog, click to choose a filter in the left column then click **Select**. You return to the **Create Contract** dialog box.

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 Network Controller

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 Network Controller GUI

This section explains how to create a cloud context profile using the Cisco Cloud Network Controller 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 13: Create Cloud Context Profile Dialog Box Fields

Properties	Description
Name	Enter the name of the cloud context profile. Match the regular expression: `[a-z]([-a-z0-9]*[a-z0-9])?` This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen.

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

Properties	Description

Add CIDR	**Note** See Understanding VPCs and Subnets Under Google Cloud and Cloud Context Profiles Under Cisco Cloud Network Controller, on page 18 for more information on primary and secondary CIDRs and subnet group labels.
	
	 To add a CIDR:
	
	 a. Click **Add CIDR**. The **Add CIDR** dialog box appears.
	
	 b. Enter the address in the **CIDR Block Range** field.
	
	 c. Click to check (enabled) or uncheck (disabled) the **Primary** check box.
	    • You must have at least one primary CIDR added for each cloud context profile.
	    • If you are adding additional secondary CIDRs and subnets for VPCs, leave the **Primary** box unchecked.
	
	 d. Click **Add Subnet** and enter the following information:
	    • In the **Address** field, enter the subnet address.
	    • In the **Name** field, enter the name for this subnet.
	    • In the **Subnet Group Label** field, choose one of the following:
	        • **Select Existing**: Click **Select Subnet Group Label**, then choose an existing subnet group label to associate with this subnet.
	        • **Create New**: Enter a unique name for the subnet group label to associate with this subnet.
	
	 e. In the **VRF** field, make a selection, if necessary.
	    • If you checked the box next to the **Primary** field, this CIDR is automatically associated with the primary VRF.
	    • If you did not check the box next to the **Primary** field, you can associate this CIDR with a secondary VRF. Click the **X** next to the VRF, then click on **Select VRF** to select the secondary VRF to associate with this CIDR.
	
	 f. When finished, click **Add**.

Step 5   Click **Save** when finished.

Configuring Virtual Machines in Google Cloud

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

This topic provides the requirements for configuring the virtual machines in Google Cloud. You can use these requirements to configure the virtual machines in Google Cloud either before you configure the endpoint selectors for Cisco Cloud Network Controller or afterward.

For example, assume that you are using **Custom** as the type of endpoint selector, as described in Endpoints and Endpoint Selectors, on page 16.

- You might go to your account in Google Cloud and create a custom tag or label in Google Cloud first, then create an endpoint selector using a custom tag or label in Cisco Cloud Network Controller afterward.

- Or you might create an endpoint selector using a custom tag or label in Cisco Cloud Network Controller first, then go to your account in Google Cloud and create a custom tag or label in Google Cloud afterward.

Before you begin

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

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

- VRF name

- Subnet information

- Google Cloud Project 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 Network Controller are displayed.

c)   Select the cloud context profile that you will use as part of this Google Cloud 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 Google Cloud virtual machine.

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

Note   For information about how to create the VM in the Google Cloud portal, see the Google Cloud documentation.

Creating a Backup Configuration Using the Cisco Cloud Network Controller 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 14: Create Backup Configuration Dialog Box Fields

Properties	Description
General	
Name	Enter the name of the backup configuration.
Description	Enter a description of the backup configuration.
Settings	

Backup Destination	Choose a backup destination.
	
	    • **Local**
	
	    • **Remote**

Properties	Description

Backup Object	Choose the root hierarchical content to consider for the backup
	
	• **Policy Universe**
	
	• **Selector Object**—When chosen, this option adds the **Object Type** drop-down list and **Object DN** field.
	
	  a. From the **Object Type** drop-down list, choose from the following options:
	
	    • **Tenant**—When chosen the **Select Tenant** option appears.
	
	    • **Application Profile**—When chosen the **Select Application Profile** option appears.
	
	    • **EPG**—When chosen the **Select EPG** option appears.
	
	    • **Contract**—When chosen the **Select Contract** option appears.
	
	    • **Filter**—When chosen the **Select Filter** option appears.
	
	    • **VRF**—When chosen the **Select VRF** option appears.
	
	    • **Cloud Context Profile**—When chosen the **Select Cloud Context Profile** option appears.
	
	  b. Click the **Select <object_name>**. The **Select <object_name>** dialog appears.
	
	  c. From the **Select <object_name>** dialog, click to choose from the options in the left column then click **Select**. You return to the **Create Backup Configuration** dialog box.
	
	    Note The **Object DN** field is automatically populated with the DN of the object it will use as root of the object tree to backup
	
	• **Enter DN**—When chosen, this option displays the **Object DN** field.
	
	  a. From the **Object DN** field, enter the DN of a specific object to use as the root of the object tree to backup.

Properties	Description
Scheduler	a. Click **Select Scheduler** to open the **Select Scheduler** dialog and choose a scheduler from the left-side column. b. Click the **Select** button at the bottom-right corner when finished.
Trigger Backup After Creation	Choose one of the following: • **Yes**—(Default) Trigger a backup after creating the backup configuration. • **No**—Do not trigger a backup after creating the backup configuration.

Step 5   Click **Save** when finished.

Creating a Tech Support Policy Using the Cisco Cloud Network Controller 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 15: Create Tech Support Dialog Box Fields

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

Properties	Description

Export Destination	Choose an export destination.
	
	   • **Controller**
	
	   • **Remote Location**—When chosen the **Select Remote Location** option appears.
	
	     a. Click **Select Remote Location**. The **Select Remote Location** dialog box appears.
	
	     b. From the **Select Remote Location** dialog, click to choose a remote location in the left column then click **Select**. You return to the **Create Tech Suport** dialog box.

Include Pre-Upgrade Logs	Click to place a check in the **Enabled** check box if you want to include pre-upgrade logs in the tech support policy.

Step 5   Click **Save** when finished.

Creating a Scheduler Using the Cisco Cloud Network Controller 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 Network Controller 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 16: Create Scheduler Dialog Box Fields

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

Properties	Description
Recurring Windows	Click **Add Recurring Window**. The **Add Recurring Window** dialog appears. a. From the **Schedule** drop-down list, choose from the following. • **Every Day** • **Even Days** • **Odd Days** • **Monday** • **Tuesday** • **Wednesday** • **Thursday** • **Friday** • **Saturday** • **Sunday** b. From the **Start Time** field, enter a time. c. In the **Maximum Concurrent Tasks** field, choose one of the following: • **Unlimited**: There is no maximum number of concurrent tasks that can be enforced on the scheduler window. • **Custom**: In the second **Maximum Concurrent Tasks** field, enter the maximum number of tasks that can be processed concurrently. The maximum value allowed in this field is 65535. d. In the **Maximum Running Time** field, choose one of the following: • **Unlimited**: There is no time limit enforced on the scheduler window. • **Custom**: In the second **Maximum Running Time** field, enter the maximum duration of the window. The acceptable format for this field is `dd:hh:mm:ss.` e. Click **Add** when finished.

Properties	Description
Add One Time Window	Click **Add One Time Window**. The **Add One Time Window** dialog appears. a. From the **Start Time** field, enter a date and time. b. From the **Maximum Concurrent Tasks** field, enter a number or leave the field blank to specify unlimited. c. From the **Maximum Running Time**, click to choose **Unlimited** or **Custom**. d. Click **Add** when finished.

Step 5 Click **Save** when finished.

Creating a Remote Location Using the Cisco Cloud Network Controller GUI

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

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 17: Create Remote Location Dialog Box Fields

Properties	Description
General	
Name	Enter the name of the remote location policy.
Description	Enter a description of the remote location policy.
Settings	
Hostname/IP Address	Enter the hostname or IP address of the remote location
Protocol	Choose a protocol: • **FTP** • **SFTP** • **SCP**

Properties	Description
Path	Enter the path for the remote location.
Port	Enter the port for the remote location.
Username	Enter a username for the remote location.
Authentication Type	When using SFTP or SCP, choose the authentication type: • **Password** • **SSH Key**
SSH Key Content	Enter the SSH key content.
SSH Key Passphrase	SSH key passphrase.
Password	Enter a password for accessing the remote location.
Confirm Password	Reenter the password for accessing the remote location.

Step 5   Click **Save** when finished.

Creating a Login Domain Using the Cisco Cloud Network Controller GUI

This section explains how to create a login domain using the Cisco Cloud Network Controller 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 18: Create Login Domain Dialog Box Fields

Properties	Description
Name	Enter the name of the login domain.
Description	Enter a description of the login domain.

Properties	Description

Realm	Choose a realm:
	
	    • **Local**
	
	    • **LDAP**—Requires adding providers and choosing an authenication type.
	
	    • **RADIUS**—Requires adding providers.
	
	    • **TACACS+**—Requires adding providers.
	
	    • **SAML**—Requires adding providers.

Providers	To add a provider: a. Click **Add Providers**. The **Select Providers** dialog appears with a list of providers in the left pane. b. Click to choose a provider. c. Click **Select** to add the provider.
Advanced Settings	Displays the **Authentication Type** and **LDAP Group Map Rules** fields.

Authentication Type	When LDAP is chosen for realm option, choose one of the following authentication types:
	
	    • **Cisco AV Pairs**—(Default)
	
	    • **LDAP Group Map Rules**—Requires adding LDAP group map rules.

Properties	Description
LDAP Group Map Rules	To add an LDAP group map rule:
	a. Click **Add LDAP Group Map Rule**. The **Add LDAP Group Map Rule** dialog appears with a list of providers in the left pane.
	b. Enter a name for the rule in the **Name** field.
	c. Enter a description for the rule in the **Description** field.
	d. Enter a group DN for the rule in the **Group DN** field.
	e. Add security domains:
	1. Click **Add Security Domain**. The **Add Security Domain** dialog box appears.
	2. Click **Select Security Domain**. The **Select Security Domain** dialog box appears with a list of security domains in the left pane.
	3. Click to choose a security domain.
	4. Click **Select** to add the security domain. You return to the **Add Security Domain** dialog box.
	5. Add a user role:
	a. From the **Add Security Domain** dialog box, click **Select Role**. The **Select Role** dialog box appears with a list of roles in the left pane.
	b. Click to choose a role.
	c. Click **Select** to add the role. You retun to the **Add Security Domain** dialog box.
	d. From the **Add Security Domain** dialog box, click the **Privilege Type** drop-down list and choose **Read Privilege** or **Write Privilege**.
	e. Click the check mark on the right side of the **Privilege Type** drop-down list to confirm.
	f. Click **Add** when finished. You return to the **Add LDAP Group Map Rule** dialog box where you can add another security domain.

Step 5   Click **Save** when finished.

Creating a Security Domain Using the Cisco Cloud Network Controller 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 **Security** > **Security Domains** > **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   In the **Type** field, choose the type of security domain:

- **Unrestricted**: Users who are assigned to this domain are able to see policies, profiles, or users configured in other security domains.

- **Restricted**: Users who are assigned to this domain will not be able to see policies, profiles, or users configured in other security domains.

Step 7   Click **Save** when finished.

Creating a Role Using the Cisco Cloud Network Controller GUI

This section explains how to create a role using the Cisco Cloud Network Controller 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 19: Create Role Dialog Box Fields

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

Properties	Description
Privilege	

Properties	Description
	Click to place a check mark in the check boxes of the privileges you want to assign the user. The privileges are:
	• **aaa**—Used for configuring authentication, authorization, accouting and import/export policies.
	• **access-connectivity**—Used for Layer 1-3 configuration under infra, static route configurations under a tenant's L3Out, management infra policies, and tenant ERSPAN policies.
	• **access-equipment**—Used for access port configuration.
	• **access-protocol**—Used for Layer 1-3 protocol configurations under infra, fabric-wide policies for NTP, SNMP, DNS, and image management, and operations-related access policies such as cluster policy and firmware policies.
	• **access-qos**—Used for changing CoPP and QoS-related policies.
	• **admin**—Complete access to everything (combine ALL roles)
	• **config-manager**
	• **custom-port-privilege**
	• **custom-privilege-1 through custom-privilege-22**
	• **fabric-connectivity**—Used for Layer 1-3 configuration under the fabric, firmware and deployment policies for raising warnings for estimating policy deployment impact, and atomic counter, diagnostic, and image management policies on leaf switches and spine switches.
	• **fabric-equipment**—Used for atomic counter, diagnostic, and image management policies on leaf switches and spine switches.
	• **fabric-protocol**—Used for Layer 1-3 protocol configurations under the fabric, fabric-wide policies for NTP, SNMP, DNS, and image management, ERSPAN and health score policies, and firmware management traceroute and endpoint tracking policies.
	• **none**—No privilege.
	• **nw-svc-params**—Used for managing Layer 4 to Layer 7 service policies.
	• **nw-svc-policy**—Used for managing Layer 4 to Layer 7 service devices and network service orchestration.
	• **ops**—Used for operational policies including monitoring and troubleshooting policies such as atomic counter, SPAN, TSW, tech support, traceroute, analytics, and core policies.
	• **site-admin**
	• **site-policy**
	• **tenant-connectivity**—Used for Layer 1-3 connectivity changes, including bridge domains, subnets, and VRFs; for atomic counter, diagnostic, and image management policies on leaf switches and spine switches; tenant in-band and out-of-band management

Properties	Description
	connectivity configurations; and debugging/monitoring policies such as atomic counters and health score.
	• **tenant-epg**—Used for managing tenant configurations such as deleting/creating endpoint groups, VRFs, and bridge domains.
	• **tenant-ext-connectivity**—Used for write access firmware policies; managing tenant L2Out and L3Out configurations; and debugging/monitoring/observer policies.
	• **tenant-ext-protocol**—Used for managing tenant external Layer 1-3 protocols, including BGP, OSPF, PIM, and IGMP, and for debugging/monitoring/observer policies such as traceroute, ping, oam, and eptrk. Generally only used for write access for firmware policies.
	• **tenant-network-profile**—Used for managing tenant configurations, such as deleting and creating network profiles, and deleting and creating endpoint groups.
	• **tenant-protocol**—Used for managing configurations for Layer 1-3 protocols under a tenant, for tenant traceroute policies, and as write access for firmware policies.
	• **tenant-qos**—Only used as Write access for firmware policies.
	• **tenant-security**—Used for Contract related configurations for a tenant.
	• **vmm-policy**—Used for managing policies for VM networking.

Step 5   Click **Save** when finished.

Creating a Certificate Authority Using the Cisco Cloud Network Controller 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 20: Create Certificate Authority Dialog Box Fields

Properties	Description
Name	Enter the name of the certificate authority.
Description	Enter a description of the certificate authority.
Used for	Choose from the following options: • **Tenant**—Choose if the certificate authority is for a specific tenant. When chosen, the **Select Tenant** option appears in the GUI. • **System**—Choose if the certificate authority is for the system.
Select Tenant	To choose a tenant: 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 Certificate Authority** dialog box.

Certificate Chain	Enter the certificate chain in the **Certificate Chain** text box.
	
	Note Add the certificates for a chain in the following order:
	
	  a. CA
	
	  b. Sub-CA
	
	  c. Subsub-CA
	
	  d. Server

Step 5   Click **Save** when finished.

Creating a Key Ring Using the Cisco Cloud Network Controller GUI

This section explains how to create a key ring using the Cisco Cloud Network Controller 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 21: Create Key Ring Dialog Box Fields

Properties	Description
Name	Enter the name of the key ring.
Description	Enter a description of the key ring.
Used for	• **System**—The key ring is for the system. • **Tenant**—The key ring is for a specific tenant. Displays a **Tenant** field for specifying the tenant.
Select Tenant	To choose a tenant: a. Click **Select Tenant**. The **Select Tenant** dialog box appears. b. From the **Select Tenant** dialog, click to choose a tenant in the left column then click **Select**. You return to the **Create Key Ring** dialog box.
Settings	
Certificate Authority	To choose a certificate authority: a. Click **Select Certificate Authority**. The **Select Certificate Authority** dialog appears. b. Click to choose a certificate authority in the column on the left. c. Click **Select**. You return to the **Create Key Ring** dialog box.
Private Key	Choose one of the following: • **Generate New Key**—Generates a new key. • **Import Existing Key**—Displays the **Private Key** text box and enables you to use an existing key.
Private Key	Enter an existing key in the **Private Key** text box (for the **Import Existing Key** option).

Properties	Description
Modulus	Click the **Modulus** drop-down list to choose from the following: • **MOD 512** • **MOD 1024** • **MOD 1536** • **MOD 2048**—(Default)
Certificate	Enter the certificate information in the **Certificate** text box.

Step 5   Click **Save** when finished.

Creating a Local User Using the Cisco Cloud Network Controller GUI

This section explains how to create a local user using the Cisco Cloud Network Controller 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 22: Create Local User Dialog Box Fields

Properties	Description
Username	Enter the username of the local user.
Password	Enter the password for the local user.
Confirm Password	Reenter the password for the local user.
Description	Enter a description of the local user.
Settings	
Account Status	To choose the account status: • **Active**—Activates the local user account. • **Blocked**—Blocks the local user account. • **Inactive**—Deactivates the local user account.

Properties	Description
First Name	Enter the first name of the local user.
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	To add a security domain:
	
	a. Click **Add Security Domain**. The **Add Security Domain** dialog box appears.
	
	b. Click **Select Security Domain**. The **Select Security Domain** dialog box appears with a list of security domains in the left pane.
	
	c. Click to choose a security domain.
	
	d. Click **Select** to add the security domain. You return to the **Add Security Domain** dialog box.
	
	e. Add a user role:
	
	  1. From the **Add Security Domain** dialog box, click **Select Role**. The **Select Role** dialog box appears with a list of roles in the left pane.
	
	  2. Click to choose a role.
	
	  3. Click **Select** to add the the role. You retun to the **Add Security Domain** dialog box.
	
	  4. From the **Add Security Domain** dialog box, click the **Privilege Type** drop-down list and choose **Read Privilege** or **Write Privilege**.
	
	  5. Click the check mark on the right side of the **Privilege Type** drop-down list to confirm.
	
	  6. Click **Add** when finished. You return to the **Create Local User** dialog box where you can add another security domain.

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

Table 23: Create Local User Dialog Box Fields: Advanced Settings

Property	Description
Account Expires	If you choose **Yes**, the account is set to expire at the time that you choose.

Property	Description
Password Update Required	If you choose **Yes**, the user must change the password upon the next login.
OTP	Put a check in the box to enable the one-time password feature for the user.
User Certificate Attribute	The attribute for the user certificate.
User Certificates	To add a user certificate: a. Click **Add X509 Certificate**. The **Add X509 Certificate** dialog box appears. b. Enter a name in the **Name** field. c. Enter the X509 certificate in the **User X509 Certificate** text box. d. Click **Add**. The **X509 certificate in the User X509 Certificate** dialog box closes. You return to the **Local User** dialog box.
SSH Keys	To add a an SSH key: a. Click **Add SSH Key**. The **Add SSH Key** dialog box appears. b. Enter a name in the **Name** field. c. Enter the SSH key in the **Key** text box. d. Click **Add**. The **Add SSH Key** dialog box closes. You return to the **Local User** dialog box.

Step 6   Click **Save** when finished.

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

With Google Cloud, the VPC resource is a global resource, which means that it spans all Google Cloud regions. By default, all regions are managed by Google Cloud and inter-region connectivity is present. Cisco Cloud Network Controller manages all 25 Google Cloud regions.

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

Step 2   In the **Workflows** area, click **Cisco Cloud Network Controller Setup**.
The **Set up - Overview** dialog box appears with options for **DNS and NTP Servers**, **Region Management**, and **Smart Licensing**.

Step 3   For **Region Management**, click **Edit Configuration**.
The **Region Management** window appears.

Step 4   Determine if you want to configure external connectivity.

Click the box next to **Enable** to enable external connectivity.

Step 5   Verify that all of the regions in the page are selected.

This page shows all of the regions that are supported by Google Cloud. All of the regions are managed by Cisco Cloud Network Controller.

Step 6   Click **Next** at the bottom of the page.

If you enabled external connectivity, the **General Connectivity** page appears.

Step 7   Enter the necessary information in the **Hub Network** area.

Hub network management is used to deploy cloud routers on specific managed regions. Configure the fabric infra connectivity for the cloud site and define the configuration template used for the cloud routers in the cloud site in this area.

Note the following restrictions:

- You can create only one hub network in Google Cloud.

- Under the hub network, only one cloud router is created in Google Cloud.

    a) In the **Hub Network** area, click **Add Hub Network**.

       The **Add Hub Network** window appears.

    b) In the **Name** field, enter a name for the hub network.
    c) Enter a value in the **BGP Autonomous System Number** field.

       The BGP Autonomous System Number (ASN) is used for BGP peering inside the cloud site and for MP-BGP IPv4 peering to other sites.

       The ASN must be a private ASN. Enter a value between 64512 and 65534 or between 4200000000 and 4294967294, inclusive, for each hub network, then click the check mark next to the field.

    d) In the **Region** field, select the appropriate regions.

       You can add up to four regions to deploy hub network in this area. The hub network will create one cloud router in each region selected.

    e) In the **VPN Router** field, enter a name for the VPN router.

       The infra VPC uses the cloud router and VPN Gateway to create IPSec tunnels and BGP sessions to on-premises sites or other cloud sites. The spoke VPCs peer with the infra VPC to share the VPN connections to external sites.

Step 8   Enter the necessary information in the **IPSec Tunnel Subnet Pools** area.

    a) In the **IPSec Tunnel Subnet Pools** area, click **Add IPSec Tunnel Subnet Pools**.

       The **Add IPSec Tunnel Subnet Pools** window appears.

    b) Enter the subnet pool to be used for IPSec tunnels, if necessary.

       By default, a subnet pool of `169.254.0.0/16` is populated to create the IPsec tunnels. You can delete the existing subnet pool and add additional subnet pools, if necessary.

The subnets used for the **IPSec Tunnel Subnet Pools** entry must be common /30 CIDRs from the `169.254.0.0/16` block. For example, `169.254.7.0/24` and `169.254.8.0/24` would be acceptable entries for the subnet pools in this field.

Click the check mark after you have entered in the appropriate subnet pools.

Step 9 When you have entered all the necessary information on this page, click **Save and Continue** at the bottom of the page.

You are given the option to create external networks and complete external connectivity configurations, if necessary. Go to Creating an External Network Using Cloud Native Routers Using the Cisco Cloud Network Controller GUI, on page 49 for those procedures.

Configuring Cisco Cloud Network Controller Using the REST API

Creating a Tenant Using the REST API

Before you begin

Review the information provided in Understanding Google Cloud Deployments with Cisco Cloud Network Controller, on page 6 before proceeding with the procedures in this section.

Step 1 Enter the following POST to share the same credentials across multiple tenants, where you are duplicating the `cloudCredentials` object in each tenant and specifying the same Google Cloud Service Account.

Note the following:

- Tenant `T1` defines the `cloudCredentials` object that carries the private key for the Service Account.

- Both tenant `T1` and `T2` then refer to this `cloudCredentials` object through the `cloudRsCredentials` relation.

- The Service Account defined by tenant `T1` must be a member of Google Cloud Projects `project1` and `project2` in this scenario.

- The highlighted areas in the POST for tenant T2 show the credentials that are shared with the first user tenant

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

<fvTenant name="T1">
    <cloudAccount id="project1" vendor="gcp" accessType="credentials" >
        <cloudRsCredentials tDn="uni/tn-T1/credentials-creds1" />
    </cloudAccount>
    <cloudCredentials name="creds1" keyId="de22a1bc-7872-4651-9d09-c5d820af7e1c"
rsaPrivateKey="-----BEGIN .... -----END PRIVATE KEY-----\n" clientId="28763876"
email="cnc-395@project2.iam.gserviceaccount.com"/>
    <fvRsCloudAccount tDn="uni/tn-T1/acct-[project1]-vendor-gcp"/>
</fvTenant>

<fvTenant name="T2">
```

```
        <cloudAccount id="project2" vendor="gcp" accessType="credentials" >
            <cloudRsCredentials tDn="uni/tn-T2/credentials-creds1" />
        </cloudAccount>
        <cloudCredentials name="creds1" keyId="de22a1bc-7872-4651-9d09-c5d820af7e1c"
rsaPrivateKey="-----BEGIN .... -----END PRIVATE KEY-----\n" clientId="28763876"
email="cnc-395@project2.iam.gserviceaccount.com"/>
        <fvRsCloudAccount tDn="uni/tn-T2/acct-[project2]-vendor-gcp"/>
</fvTenant>
```

Step 2 To create a user tenant where the Cisco Cloud Network Controller runs outside of Google Cloud (the infra tenant with credentials):

Note that the new properties added specifically for Google Cloud are highlighted below.

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

<fvTenant name="infra">
    <cloudAccount id="project1" vendor="gcp" accessType="credentials" >
        <cloudRsCredentials tDn="uni/tn-infra/credentials-creds1" />
    </cloudAccount>
    <cloudCredentials name="creds1" keyId="de22a1bc-7872-4651-9d09-c5d820af7e1c"
rsaPrivateKey="-----BEGIN .... -----END PRIVATE KEY-----\n" clientId="28763876"
email="cnc-395@project2.iam.gserviceaccount.com"/>
    <fvRsCloudAccount tDn="uni/tn-infra/acct-[project1]-vendor-gcp"/>
</fvTenant>

<fvTenant name="T2">
    <cloudAccount id="project2" vendor="gcp" accessType="credentials" >
        <cloudRsCredentials tDn="uni/tn-infra/credentials-creds1" />
    </cloudAccount>
    <fvRsCloudAccount tDn="uni/tn-T2/acct-[project2]-vendor-gcp"/>
</fvTenant>
```

Step 3 To create a managed user tenant where the user shares the infra service account across multiple Google Cloud projects:

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

<fvTenant name="infra">
    <cloudAccount id="project1" vendor="gcp" accessType="managed" />
    <fvRsCloudAccount tDn="uni/tn-infra/acct-[project1]-vendor-gcp"/>
</fvTenant>

<fvTenant name="T2">
    <cloudAccount id="project2" vendor="gcp" accessType="managed" />
    <fvRsCloudAccount tDn="uni/tn-T2/acct-[project2]-vendor-gcp"/>
</fvTenant>
```

Configuring Inter-VRF Route Leaking Using the REST API

This example demonstrates how to configure leak routes for the Cisco Cloud Network Controller using the REST API. This example shows how to configure inter-VRF route leaking, between an external VRF and a cloud VRF, as shown in the following figure.
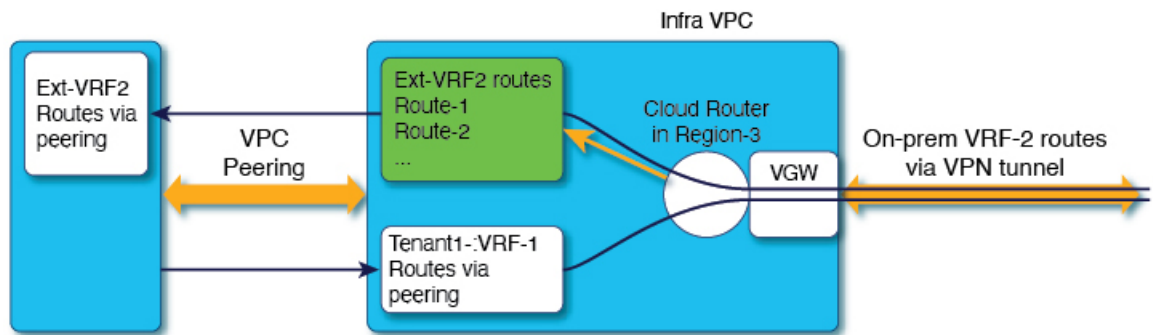
Subnet1 (Region-1) Route-Table

CIDR1 (Region-1) - 100.100.0.0/16
Subnet1 - 100.100.100.0/24

100.100.0.0/16 -> Local
50.50.0.0/16 -> Infra-VPC

Leak-All-routes to
Tenant-Infra:Ext-RF-2

To configure inter-VRF route leaking for this example:

Example:

```
<polUni>
    <fvTenant name="t1">
        <fvCtx name="VRF1">
            <leakRoutes>
                <leakInternalPrefix ip="0.0.0.0/0" status="">
                    <leakTo tenantName="infra" ctxName="Ext-VRF2" scope="public" status=""/>
                </leakInternalPrefix>
            </leakRoutes>
        </fvCtx>
        <cloudCtxProfile name="v1-us-west1" type="regular" vpcGroup="one" status="">
            <cloudRsToCtx tnFvCtxName="VRF1"/>
            <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-gcp/region-us-west1"/>
            <cloudCidr addr="100.100.0.0/16" primary="yes">
                <cloudSubnet ip="100.100.100.0/20" scope="public,shared" subnetGroup="one">
                    <cloudRsZoneAttach tDn="uni/clouddomp/provp-gcp/region-us-west1/zone-default"/>
                </cloudSubnet>
            </cloudCidr>
        </cloudCtxProfile>
    </fvTenant>
    <fvTenant name="infra" status="">
        <fvCtx name="Ext-VRF2">
            <leakRoutes>
                <leakExternalPrefix ip="0.0.0.0/0" status="">
                    <leakTo tenantName="t1" ctxName="VRF1" scope="public" status=""/>
                </leakInternalPrefix>
            </leakRoutes>
        </fvCtx>
    </fvTenant>
</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>
    <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 a Contract Using the REST API

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

Before you begin

Create filters.

To create a contract:

```
<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>
```

Note the following restrictions for the name of the contract (the `vzBrCP` entry):

- Match the regular expression:

  `[a-z]([-a-z0-9]*[a-z0-9])?`

  This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen.

- We recommend using 14 characters or fewer for this name, if possible, due to the restrictions imposed by the Google Cloud firewall rules. Refer to Naming Length Restrictions Imposed By Google Cloud Firewall Rules, on page 23 to better understand the restriction and the total number of characters allowed for each of the Cisco Cloud Network Controller components that make up the firewall rule name.

Creating a Cloud Context Profile Using the REST API

This section demonstrates how to create a cloud context profile.

Before you begin

Create a VRF.

Step 1 To create a basic cloud context profile:

Example:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
    <fvTenant name="tn15">
    <cloudCtxProfile name="cProfilewest1151">
        <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-gcp/region-us-west1"/>
      <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-gcp/region-us-west1/zone-default"/>
        </cloudSubnet>
        <cloudSubnet ip="15.151.2.0/24" name="albsubnet" >
            <cloudRsZoneAttach tDn="uni/clouddomp/provp-gcp/region-us-west1/zone-default"/>
        </cloudSubnet>
        <cloudSubnet ip="15.151.3.0/24" name="subnet" usage="">
            <cloudRsZoneAttach tDn="uni/clouddomp/provp-gcp/region-us-west1/zone-default"/>
        </cloudSubnet>
    </cloudCidr>
    </cloudCtxProfile>
    </fvTenant>
</polUni>
```

Step 2 To create a cloud context profile where you are adding a secondary VRF, CIDR, and subnet for a VNet:

Example:

```xml
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
    <fvTenant name="tenant1" status="">
        <fvCtx name="VRF1" />
        <fvCtx name="VRF2" />
        <cloudCtxProfile name="vpc1" status="">
            <cloudRsCtxProfileToRegion tDn="uni/clouddomp/provp-gcp/region-us-central1" status=""/>
            <cloudRsToCtx tnFvCtxName="VRF1" />
            <cloudRsCtxProfileToGatewayRouterP tDn="uni/tn-infra/gwrouterp-default" status=""/>
            <cloudCidr name="cidr1" addr="192.0.2.0/16" primary="yes" status="">
                <cloudSubnet ip="192.0.3.0/24" usage="gateway" status="">
                    <cloudRsZoneAttach status=""
tDn="uni/clouddomp/provp-gcp/region-us-central1/zone-default"/>
                </cloudSubnet>
            </cloudCidr>
            <cloudCidr name="cidr1" addr="193.0.2.0/16" primary="no" status="">
                <cloudSubnet ip="193.0.3.0/24" usage="" status="">
                    <cloudRsSubnetToCtx tnFvCtxName="VRF2"/>
                    <cloudRsZoneAttach status=""
tDn="uni/clouddomp/provp-gcp/region-us-central1/zone-default"/>
                </cloudSubnet>
            </cloudCidr>
        </cloudCtxProfile>
    </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-[<gcp-id>]-vendor-gcp" />


        <fvCtx name="ctx151"/>

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

    </cloudApp>

  </fvTenant>
</polUni>
```

For the application profile name, note the following restrictions:

- Match the regular expression:

  ```
  [a-z]([-a-z0-9]*[a-z0-9])?
  ```

  This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen.

- We recommend using 14 characters or fewer for this name, if possible, due to the restrictions imposed by the Google Cloud firewall rules. Refer to Naming Length Restrictions Imposed By Google Cloud Firewall Rules, on page 23 to better understand the restriction and the total number of characters allowed for each of the Cisco Cloud Network Controller components that make up the firewall rule name.

Creating an EPG Using the REST API

Use the procedures in this section to create an application EPG or an external EPG using the REST API.

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:

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


        <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>
```

Note the following restrictions:

- Match the regular expression:

```
[a-z]([-a-z0-9]*[a-z0-9])?
```

This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen.

• We recommend using 14 characters or fewer for this name, if possible, due to the restrictions imposed by the Google Cloud firewall rules. Refer to Naming Length Restrictions Imposed By Google Cloud Firewall Rules, on page 23 to better understand the restriction and the total number of characters allowed for each of the Cisco Cloud Network Controller components that make up the firewall rule name.

Creating an External Cloud EPG Using the REST API

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

For the name of the external EPG, note the following restrictions:

• Match the regular expression:

```
[a-z]([-a-z0-9]*[a-z0-9])?
```

This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen.

• We recommend using 14 characters or fewer for this name, if possible, due to the restrictions imposed by the Google Cloud firewall rules. Refer to Naming Length Restrictions Imposed By Google Cloud Firewall Rules, on page 23 to better understand the restriction and the total number of characters allowed for each of the Cisco Cloud Network Controller components that make up the firewall rule name.

Before you begin

Create an application profile and a VRF.

Step 1 To create an external cloud EPG:

Example:

```xml
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
    <fvTenant name="tn15">
        <fvRsCloudAccount tDn="uni/tn-infra/act-[<gcp-id>]-vendor-gcp" />
        <fvCtx name="ctx151"/>
        <cloudVpnGwPol name="VgwPol1"/>
        <cloudApp name="a1">
        <cloudExtEPg routeReachability="internet" name="extEpg-1">
            <fvRsCons tnVzBrCPName="extEpg-1"/>
            <cloudRsCloudEPgCtx tnFvCtxName="ctx151"/>
            <cloudExtEPSelector name="extSelector1" subnet="0.0.0.0/0"/>
        </cloudExtEPg>
    </cloudApp>
  </fvTenant>
</polUni>
```

Step 2 To create an external cloud EPG with type **site-external**, or an infra L3Out EPG:

Example:

```xml
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/.xml -->
<polUni>
    <fvTenant name="infra">
        <cloudApp name="a1">
        <cloudExtEPg routeReachability="site-ext" name="extEpg-1">
            <fvRsCons tnVzBrCPName="extEpg-1"/>
            <cloudRsCloudEPgCtx tnFvCtxName="ctx152"/>
            <cloudExtEPSelector name="extSelector1" subnet="10.100.0.0/16"/>
        </cloudExtEPg>
    </cloudApp>
  </fvTenant>
</polUni>
```

Creating Cloud Routers, External Networks, and External VRFs Using the REST API

This section demonstrates how to create cloud routers, external networks, and external VRFs using the REST API.

Following is an example POST that shows how to bring up the cloud router in four regions and add an external network with an external VRF in each region.

```xml
<polUni>
    <fvTenant name="infra" status="">
        <fvCtx name="extv1" pcEnfPref="enforced" status=""/>
         <fvCtx name="extv2" pcEnfPref="enforced" status=""/>
          <fvCtx name="extv3" pcEnfPref="enforced" status=""/>

          <cloudtemplateInfraNetwork name="default" vrfName="overlay-1" hostRouterMode="manual"
status="">
            <cloudtemplateIpSecTunnelSubnetPool subnetpool= "169.254.7.0/24" poolname="pool1" />
             <cloudtemplateIpSecTunnelSubnetPool subnetpool= "169.254.8.0/24" poolname="pool2" />
              <cloudtemplateIpSecTunnelSubnetPool subnetpool= "169.254.10.0/24" poolname="pool3" />

            <cloudtemplateHubNetwork name="default" status="" >
                <cloudtemplateHubNetworkName name="foo1" asn="64514" status="">
                    <cloudRegionName provider="gcp" region="us-west4" status="" />
                    <cloudRegionName provider="gcp" region="us-west2" status="" />
                    <cloudRegionName provider="gcp" region="us-east1" status="" />
                      <cloudRegionName provider="gcp" region="us-west1" status=""/>
                </cloudtemplateHubNetworkName>
            </cloudtemplateHubNetwork>

            <cloudtemplateIntNetwork name="default">
                <cloudRegionName provider="gcp" region="us-west1">
                    <cloudtemplateVpnRouter name="default" status=""/>
                </cloudRegionName>
                <cloudRegionName provider="gcp" region="us-west2">
                    <cloudtemplateVpnRouter name="default" status=""/>
                </cloudRegionName>
                <cloudRegionName provider="gcp" region="us-east1">
                    <cloudtemplateVpnRouter name="default" status=""/>
                  </cloudRegionName>
                    <cloudRegionName provider="gcp" region="us-west4">
                    <cloudtemplateVpnRouter name="default" status=""/>
```

```
                    </cloudRegionName>
              </cloudtemplateIntNetwork>

              <cloudtemplateExtNetwork name="default">
          </cloudtemplateExtNetwork>
                <cloudtemplateExtNetwork name="extnwfoo1" vrfName="extv1" hubNetworkName="foo1"
vpnRouterName="default" status="">
                    <cloudRegionName provider="gcp" region="us-west1" status=""/>
                        <cloudtemplateVpnNetwork name="onprem01" remoteSiteId="1" status="">
                  <cloudtemplateIpSecTunnel peeraddr="128.1.1.1" preSharedKey="abcd" poolname="pool1"
 status="">
                              <cloudtemplateBgpIpv4 peeraddr="0.0.0.0/0" peerasn="64529" status=""/>
                        </cloudtemplateIpSecTunnel>
                    </cloudtemplateVpnNetwork>
              </cloudtemplateExtNetwork>
              <cloudtemplateExtNetwork name="extnwfoo2" vrfName="extv2" hubNetworkName="foo1"
vpnRouterName="default" status="">
                 <cloudRegionName provider="gcp" region="us-west2" status=""/>
                    <cloudtemplateVpnNetwork name="onprem02" remoteSiteId="2" status="">
                        <cloudtemplateIpSecTunnel peeraddr="128.1.1.2" preSharedKey="def"
poolname="pool2" status="">
                              <cloudtemplateBgpIpv4 peeraddr="0.0.0.0/0" peerasn="64529" status=""/>
                        </cloudtemplateIpSecTunnel>
                    </cloudtemplateVpnNetwork>
              </cloudtemplateExtNetwork>
                <cloudtemplateExtNetwork name="extnwfoo3" vrfName="extv3" hubNetworkName="foo1"
vpnRouterName="default" status="">
                    <cloudRegionName provider="gcp" region="us-east1" status=""/>
                      <cloudtemplateVpnNetwork name="onprem03" remoteSiteId="3" status="">
                            <cloudtemplateIpSecTunnel peeraddr="128.1.1.3" preSharedKey="abc"
poolname="pool3" status="">
                              <cloudtemplateBgpIpv4 peeraddr="0.0.0.0/0" peerasn="64529" status=""/>
                        </cloudtemplateIpSecTunnel>
                    </cloudtemplateVpnNetwork>
                  </cloudtemplateExtNetwork>
          </cloudtemplateInfraNetwork>
      </fvTenant>
</polUni>
```

Viewing System Details

Monitoring VM Host Metrics

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

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

The Prometheus Node Exporter is automatically available by default.

Guidelines and Limitations

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

Monitoring VM Host Metrics Using the GUI

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

Step 1   In the Cisco Cloud Network Controller GUI, navigate to **Infrastructure** > **System Configuration**, then click on the **Management Access** tab.

Step 2   In the **HTTPS** area to the right of the window, note the entry in the **Node Exporter** field.
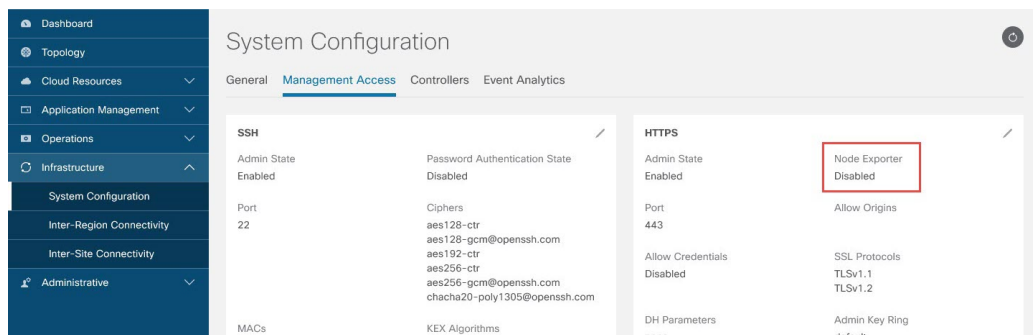
- **Enabled**: The Prometheus Node Exporter has already been enabled. You do not have to continue with these instructions in that case.

- **Disabled**: The Prometheus Node Exporter is not enabled yet. Proceed with these instructions to enable the Prometheus Node Exporter.

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

The **HTTPS Settings** window appears.

Step 4   Locate the **Node Exporter** field and click **Enable**.
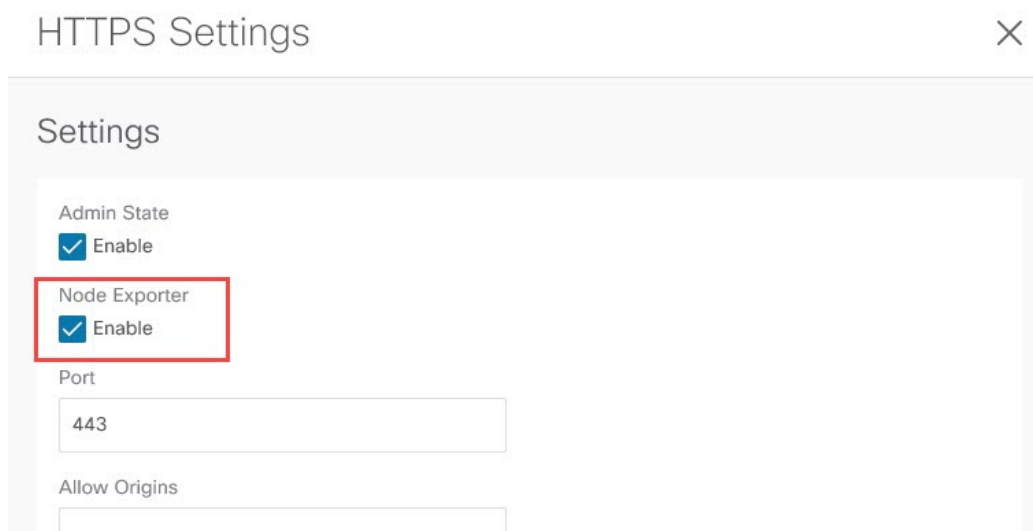
A warning message appears, telling you that saving these settings will restart the web service, and that it will take a moment for it to resume responding to requests. Click **OK** to confirm these changes.

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

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

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

This verifies that the Prometheus Node Exporter is enabled.
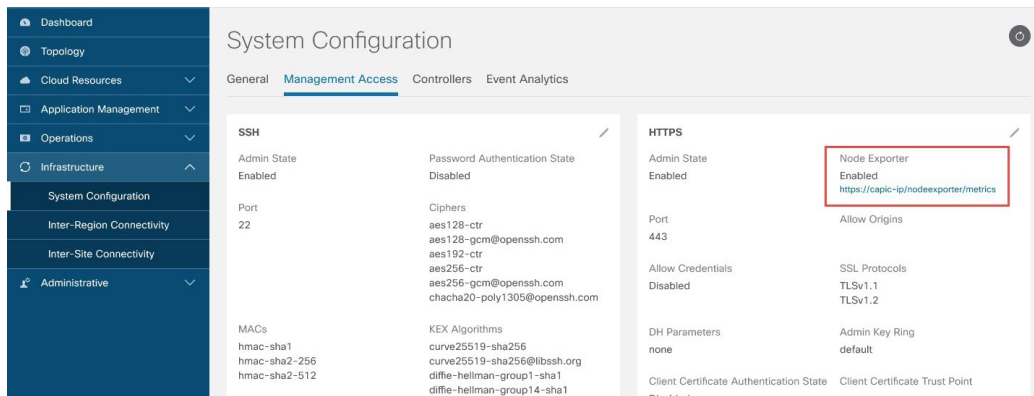
Step 7 Click the link under the **Enabled** text in the **Node Exporter** area.

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

Monitoring VM Host Metrics Using the REST API

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

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

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

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

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

```
POST https://<cloud-network-controller-ip-address>/api/mo/uni/fabric/comm-default/https.xml

<commHttps nodeExporter="enabled" />
```

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

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

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

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

```
POST https://<cloud-network-controller-ip-address>/api/mo/uni/fabric/comm-default/https.xml

<commHttps nodeExporter="disabled" />
```

Viewing Application Management Details

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

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.
Cloud Context Profiles	Displays cloud context profiles as rows in a summary table.
External Networks	Displays external networks 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.

- **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 Network Controller GUI. The cloud resource details include the information about a specific region, VPC, router, security group (application security group/network security group), endpoint, VM, and cloud service.

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

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

Table 25: Cloud Resource Subtabs

Subtab Name	Description
Regions	Displays regions as rows in a summary table.
VPCs	Displays VPCs as rows in a summary table.
Routers	Displays routers as rows in a summary table.
Endpoints	Displays endpoints as rows in a summary table.
Instances	Displays instances as rows in a summary table.

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

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

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

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, including the cloud tags associated with endpoints.

- **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.

- **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 Network Controller 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
Event Analytics	Contains the following subtabs: • **Faults** Tab—Displays faults as rows in a summary table. • **Fault Records** Tab—Displays fault records as rows in a summary table. • **Events** Tab—Displays events as rows in a summary table. • **Audit Logs** Tab—Displays audit logs as rows in a summary table.
Active Sessions	Displays a list of active users who are logged into Cisco Cloud Network Controller.

Subtab Name	Description

Backup & Restore	Contains the following subtabs:
	
	• **Backups** Tab—Displays backup as rows in a summary table.
	
	• **Backup Policies** Tab—Displays backup policies as rows in a summary table.
	
	• **Job Status** Tab—Displays the job status as rows in a summary table.
	
	• **Event Analytics** Tab—Contains the following subtabs:
	    • **Faults** Tab—Displays faults as rows in a summary table.
	    • **Events** Tab—Displays events as rows in a summary table.
	    • **Audit Logs** Tab—Displays audit logs as rows in a summary table.

Tech Support	Contains the following subtabs: • **Tech Support**Tab—Displays tech support policies as rows in a summary table. • **Core Logs** Tab—Displays core logs as rows in a summary table.

Firmware Management	Contains the following subtabs:
	
	• **Controllers** Tab—Displays general firmware management information, such as Current Firmware Version, Upgrade Status, and so on.
	
	• **Images** Tab—Displays a list of images.
	
	• **Event Analytics** Tab—Contains the following subtabs:
	    • **Faults** Tab—Displays faults as rows in a summary table.
	    • **Events** Tab—Displays events as rows in a summary table.
	    • **Audit Logs** Tab—Displays audit logs as rows in a summary table.

Schedulers	Displays scheduler policies as rows in a summary table.
Remote Locations	Displays remote locations as rows in a summary table.

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

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

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

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 Network Controller 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
System Configuration	Displays **General** system configuration information, **Management Access** information, **Controllers**, and **Event Analytics**.
External Connectivity	Displays one pane with a map that contains the inter-region connectivity view.

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 Network Controller 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
Authentication	Displays the **Authentication Default Settings**, **Login Domains**, **Providers** and **Event Analytics** subtabs, which contain the information described below: • **Authentication Default Settings** Tab—Displays settings information. • **Login Domains** Tab—Displays the login domains as rows in a summary table. • **Providers** Tab—Displays the providers as rows in a summary table. • **Event Analytics** Tab—Displays the **Faults**, **Events**, and **Audit Logs** subtabs, each with the corresponding information displayed as rows in a summary table.
Security	Contains the following list of subtabs: • **Security Default Settings** Tab—Enables you to view the default security settings information. • **Security Domains** Tab—Enables you to view security domain information in a summary table. • **Roles** Tab—Enables you to view the role information in a summary table. • **RBAC Rules** Tab—Enables you to view RBAC rule information in a summary table. • **Certificate Authorities** Tab—Enables you to view the certificate authority information in a summary table. • **Key Rings** Tab—Enables you to view key ring information in a summary table. • **User Activity** Tab—Enables you to view user activity.
Users	Contains the following subtabs: • **Local** Tab—Displays local users as rows in a summary table. • **Remote** Tab—Displays remote users as rows in a summary table.

Subtab Name	Description
Smart Licensing	Contains the following subtabs: • **General** Tab—Displays the licenses as rows in a summary table. • **CSRs** Tab—Displays CSRs as rows in a summary table. • **Faults** Tab—Displays faults as rows in a summary table.

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

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

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 Network Controller GUI

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

• Regions

• VPCs

• Endpoints

• Instances

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

The **Dashboard** window for the Cisco Cloud Network Controller system appears. From this window, you can view the overall health status of your system.
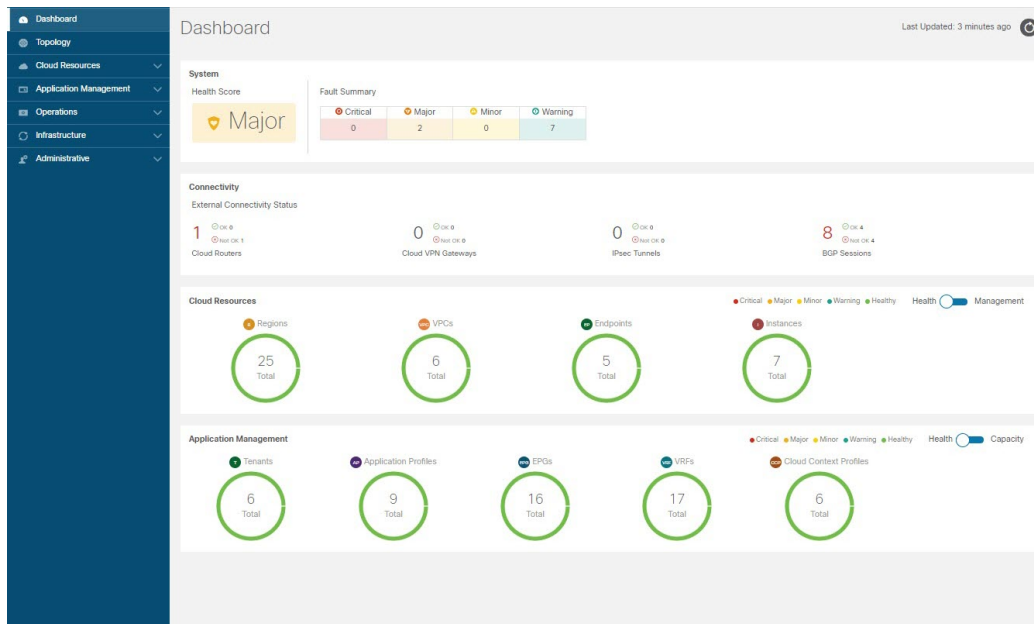
Step 2   Click within the Fault Summary area in the **Dashboard** window.

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.

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.

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**. The health of each region is displayed in the left column of the table in the **Regions** window.

Cisco Cloud Network Controller Statistics

About Google Cloud Statistics

Beginning in Cisco Cloud Network Controller Release 25.0(4), you can view statistics that are derived by processing Google Cloud flow logs.

Beginning in Cisco Cloud Network Controller Release 25.0(5), you can view statistics that are collected from the Cisco Catalyst 8000V routers.

Router Statistics

Cisco Cloud Network Controller allows you to view router statistics for individual cloud context profiles within a tenant. Statistics are displayed for Cisco Catalyst 8000V routers within the cloud context profile.

Each router instance captures and stores the ingress and egress byte and packet statistics for each physical and tunnel interface. The Cisco Cloud Network Controller queries the routers for these statistics and maps the response to router statistics on the Cisco Cloud Network Controller. The statistics query repeats every 5 minutes for as long as the tunnel is up and operational. The Cisco Cloud Network Controller GUI displays the collected statistics for the routers.

Flow Log Statistics

Cisco Cloud Network Controller allows you to enable flow log statistics for individual cloud context profiles within a tenant. When statistics are enabled for a cloud context profile, statistics are collected for every IP address within the corresponding VPCs. Available statistics include ingress and egress bytes and packets, internal and external, for VPCs, regions, and endpoints.

The collected statistics are aggregated through the following hierarchy:

• IP statistics are aggregated to determine endpoint statistics.

• Endpoint statistics are aggregated to determine zone statistics.

• Zone statistics are aggregated to determine subnet statistics.

• Subnet statistics are aggregated to determine region statistics.

• Region statistics are aggregated to determine VPC statistics.

The Cisco Cloud Network Controller GUI displays the collected statistics for VPCs, regions, and endpoints.

For more information about Google Cloud flow logs, see "VPC Flow Logs" on the Google Cloud website.

Guidelines and Limitations For Configuring Google Cloud Statistics

Following are the guidelines and limitations when configuring Cisco Cloud Network Controller to collect Google Cloud statistics:

• Router statistics are enabled by default when Cisco Catalyst 8000V routers are brought up in the Cisco Cloud Network Controller.

• The flow log statistics feature is not enabled by default.

• Flow log statistics can be enabled for individual context profiles within a tenant. In this case, flow logs are enabled on all subnets belonging to the corresponding VPC.

• Flow logs are aggregated at one minute intervals. The aggregation interval and sample rate are not configurable.

• Statistics for dropped traffic are not supported by flow logs.

• Statistics filters are not supported.

• Zone and subnet statistics are not displayed.

Viewing Router Statistics

In the Cisco Cloud Network Controller GUI, you can view graphed statistics for Cisco Catalyst 8000V routers. Available statistics for each router include ingress and egress bytes and packets. Peak values for each counter are displayed with a timestamp that shows when the peak value occurred.

This example procedure shows you how to view the statistics for a router in the Cisco Cloud Network Controller GUI.

Step 1   Select the Cisco Catalyst 8000V router whose statistics you would like to view.

a) From the navigation menu, select **Cloud Resources** > **Routers**.

A **Routers** summary table appears in the work pane.

b) In the summary table, double-click the name of the router.

The router dialog box appears over the work pane. The router dialog box displays the **Overview**, **Cloud Resources**, **Application Management**, **Statistics**, and **Event Analytics** tabs.

Step 2   Click the **Statistics** tab.
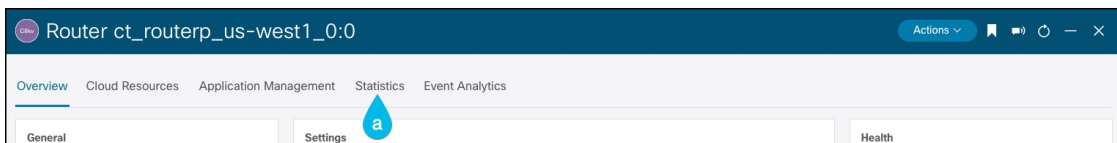
A graphical view of the router statistics appears along with a table of maximum values.

Step 3   Configure the display of the statistics.

To configure the parameters of the displayed statistics, you can modify the following settings:

Properties	Description
Sampling Interval	Choose the interval: • 1 hour • 12 Hours • 1 Day • 1 Week • 1 Month
Stats Type	Choose the display type: • Periodic • Cumulative • Trend • Rate
Bytes	Select the checkbox to display the byte counter graph. The vertical axis on the left side of the graph indicates the byte count.
Packets	Select the checkbox to display the packet counter graph. The vertical axis on the right side of the graph indicates the packet count.

Enabling Flow Log Statistics

You can enable the collection of Google Cloud flow log statistics for individual context profiles within a tenant. Statistics can then be viewed for VPCs, regions, and endpoints in their respective **Cloud Resources** GUI menus.

To enable flow log statistics using the Cisco Cloud Network Controller GUI:

Step 1   Select the tenant containing the resource for which flow log statistics will be enabled.
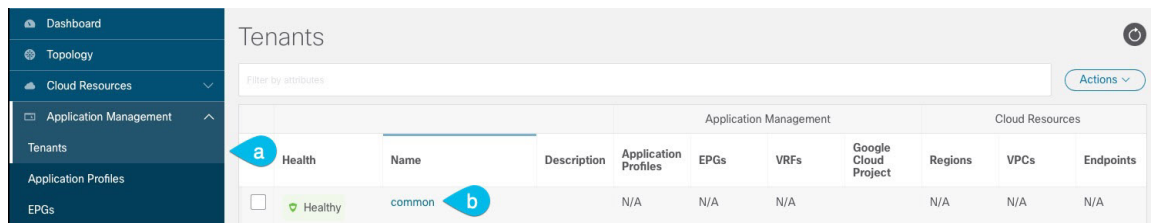
a) From the navigation menu, select **Application Management** > **Tenants**.
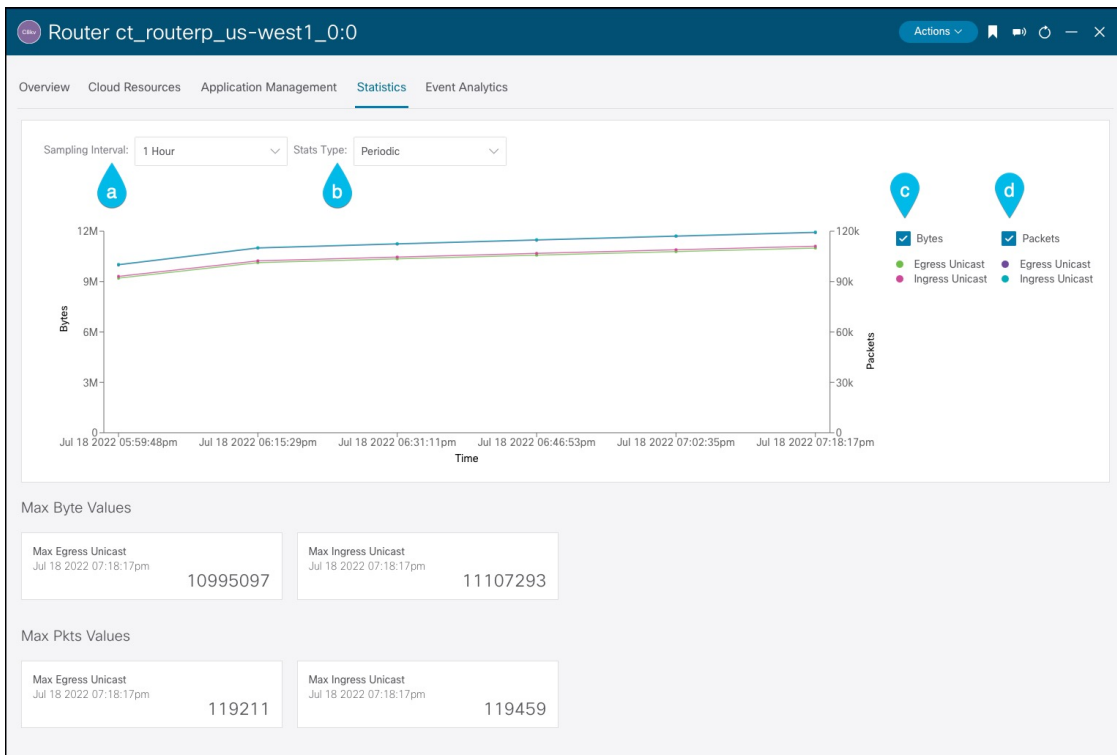
A **Tenants** summary table appears in the work pane.

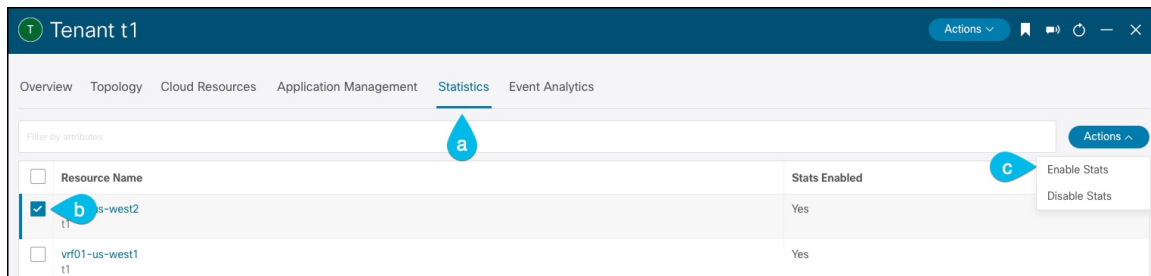b) In the summary table, double-click the name of the tenant.

The tenant dialog box appears over the work pane. The tenant dialog box displays the **Overview**, **Topology**, **Cloud Resources**, **Application Management**, **Statistics**, and **Event Analytics** tabs.

Step 2 Enable flow log statistics collection on the desired resource.

a) In the tenant dialog box, select the **Statistics** tab.

A **Resource Name** table appears with context profiles listed as rows in the table. The **Stats Enabled** column indicates whether flow log statistics are enabled for each resource.

b) Check the checkbox next to the desired resource.
c) In the top right of the tenant dialog box, click the **Actions** menu and select **Enable Stats**.

Viewing Flow Log Statistics

In the Cisco Cloud Network Controller GUI, you can view graphed statistics for VPCs, regions, and endpoints. Available statistics for each include ingress and egress bytes and packets. For VPCs and regions, the statistics are further separated into the following categories:

- Statistics: All traffic counters extracted from the flow logs records.

- Inter-Region Statistics: For a particular region, all ingress and egress traffic to or from other regions within the VPC.

- External Statistics: All ingress and egress traffic with a source or destination outside of the VPC.

- Inter-Zone Statistics: For a particular zone, all ingress and egress traffic to or from other zones within the same region and VPC. These aggregated statistics are available on the region page and on the VPC page.

Peak values for each counter are displayed with a timestamp that shows when the peak value occurred.

This example procedure shows you how to view the flow log statistics for VPCs in the Cisco Cloud Network Controller GUI. You can also view the statistics for regions or endpoints in the same manner by selecting **Regions** or **Endpoints** instead of **VPCs** in the following steps.

Before you begin

Enable Google Cloud flow log statistics for the desired cloud context profile using the procedure in Enabling Flow Log Statistics, on page 126.

Step 1   Select the resource whose statistics you would like to view.
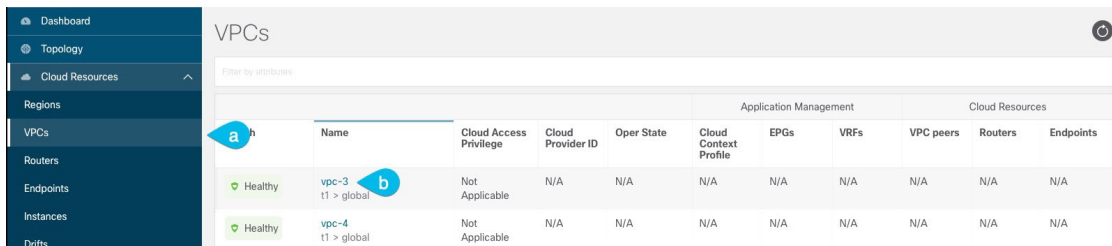
a)   From the navigation menu, select **Cloud Resources** > **VPCs**.

A **VPCs** summary table appears in the work pane.

b)   In the summary table, double-click the name of the VPC.

The VPC dialog box appears over the work pane. The VPC dialog box displays the **Overview**, **Application Management**, **Cloud Resources**, **Statistics**, and **Event Analytics** tabs.

Step 2   Click the **Statistics** tab.
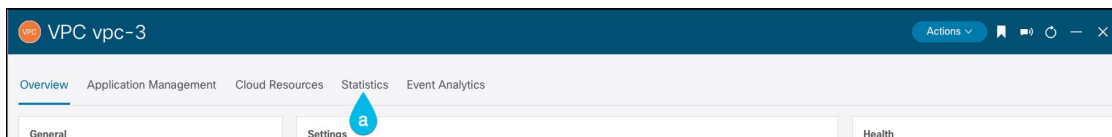
A graphical view of the VPC statistics appears along with a table of maximum values.

Step 3 Configure the display of the statistics.

To configure the parameters of the displayed statistics, you can modify the following settings:

Properties	Description
Sampling Interval	Choose the interval: • 1 hour • 12 Hours • 1 Day • 1 Week • 1 Month
Stats Type	Choose the display type: • Periodic • Cumulative • Trend • Rate

Properties	Description
Bytes	Check the checkbox to display the byte counter graph. The vertical axis on the left side of the graph indicates the byte count.
Packets	Check the checkbox to display the packet counter graph. The vertical axis on the right side of the graph indicates the packet count.

Enabling VPC Flow Log Statistics Using the REST API

Google Cloud flow log statistics can be enabled for individual context profiles within a tenant.

Step 1   Define a flow log policy (`cloudGcpFlowLogPol`) under the tenant.

No configuration settings are needed except for the name.

Note   For the name of the flow log policy, note the following restrictions:

- Match the regular expression:

  `[a-z]([-a-z0-9]*[a-z0-9])?`

  This means that the first character must be a lowercase letter, and all the following characters must be hyphens, lowercase letters, or digits, except the last character, which cannot be a hyphen.

- We recommend using 14 characters or fewer for this name.

Example:

```
<polUni>
    <fvTenant name="tenant1" status="">
        <cloudGcpFlowLogPol name="myFlowLogPol1" status="">
        </cloudGcpFlowLogPol>
        <cloudCtxProfile name="ctxProfile2" status="" vpcGroup="vpc-4">
.
.
.
```

Step 2   Within the cloud context profile, add a reference to the flow log policy.

Flow log statistics for the cloud context profile are enabled by the presence of the reference object (`cloudRsCtxToGcpFlowLog`). To disable flow log statistics for the cloud context profile, remove the reference object.

Example:

```
.
.
.

            <cloudRsCtxToGcpFlowLog tnCloudGcpFlowLogPolName="myFlowLogPol1" status=""/>
        </cloudCtxProfile>
    </fvTenant>
```

```
</polUni>
```

Cisco Cloud Network Controller Security

This chapter contains the following sections:

Access, Authentication, and Accounting

Cisco Cloud Application Policy Infrastructure Controller (Cisco Cloud Network Controller) 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 Network Controller GUI, on page 96 to configure a Local User and associate it to the OTP, SSH Public Key, and X.509 User Certificate using the Cisco Cloud Network Controller GUI.

Configuring TACACS+, RADIUS, LDAP and SAML Access

The following topics describe how to configure TACACS+, RADIUS, LDAP and SAML access for the Cisco Cloud Network Controller.

Overview

This topic provides step-by-step instructions on how to enable access to the Cisco Cloud Network Controller 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 Cisco Cloud Network Controller for TACACS+ Access

Before you begin

- The Cloud Application Policy Infrastructure Controller (Cisco Cloud Network Controller) is online.

- The TACACS+ server host name or IP address, port, and key are available.

- The Cisco Cloud Network Controller management endpoint group is available.

Step 1 In the Cisco Cloud Network Controller, create the **TACACS+ 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 **TACACS+**.

f) In **Settings** section, specify the **Key** and **Confirm Key**, **Port**, **Authentication Protocol**, **Timeout**, **Retries**, **Management EPG**. Select either **Enabled** or **Disabled** for **Server Monitoring**.

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

a) On the menu bar, choose **Administrative** > **Authentication**.

b) In the Work pane, click on **Login Domains** tab and then click on the **Actions** drop-down and select **Create Login Domain**.

The **Create Login Domain** dialog box appears.

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

Properties	Description
General	

Properties	Description
Name	Enter the name of the Login Domain
Description	Enter the description of the Login Domain.
Settings	
Realm	Choose **TACACS**+ from the dropdown menu
Providers	To choose a Provider(s): 1. Click **Add Providers**. The **Select Providers** dialog appears. 2. Click to choose a provider(s) in the column on the left. 3. Click **Select**. You return to the **Create Login Domain** dialog box.

d) Click **Save** to save the configuration.

What to do next

This completes the TACACS+ configuration steps. Next, if a RADIUS server will also be used, configure the Cisco Cloud Network Controller for RADIUS.

Configuring Cisco Cloud Network Controller for RADIUS Access

Before you begin

- The Cloud Application Policy Infrastructure Controller (Cisco Cloud Network Controller) is online.

- The RADIUS server host name or IP address, port, and key are available.

- The Cisco Cloud Network Controller management endpoint group is available.

Step 1 In the Cisco Cloud Network Controller, create the **RADIUS 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 **RADIUS**.

f) In the **Settings** section, specify the **Key** and **Confirm 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) On the menu bar, choose **Administrative** > **Authentication**.

b) In the Work pane, click on **Login Domains** tab and then click on the **Actions** drop-down and select **Create Login Domain**.

The **Create Login Domain** dialog box appears.

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

Properties	Description
General	
Name	Enter the name of the Login Domain
Description	Enter the description of the Login Domain.
Settings	
Realm	Choose **RADIUS** from the dropdown menu
Providers	To choose a Provider(s): 1. Click **Add Providers**. The **Select Providers** dialog appears. 2. Click to choose a provider(s) in the column on the left. 3. Click **Select**. You return to the **Create Login Domain** dialog box.

d) Click **Save** to save the configuration.

What to do next

This completes the Cisco Cloud Network Controller RADIUS configuration steps. Next, configure the RADIUS server.

Configuring a Cisco Secure Access Control Server for RADIUS and TACACS+ Access to the Cisco Cloud Network Controller

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 Cisco Cloud Network Controller

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 Cisco Cloud Network Controller for LDAP Access

Before you begin

> • The Cisco Cloud Network Controller is online.
>
> • The LDAP server host name or IP address, port, bind DN, Base DN, and password are available.
>
> • The Cisco Cloud Network Controller management endpoint group is available.

Step 1 In the Cisco Cloud Network Controller, 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**, **Confirm Password**, **Port**, **Timeout**, **Retries**, **SSL**, **SSL Certificate Validation Level**, **Attribute**, **Filter Type**, **Management EPG**, and **Server Monitoring**.

In the **SSL Certificate Validation Level** field, you have the following options:

> • **Permissive**: A debugging knob to help diagnose DUO LDAP SSL Certificate issues.
>
> • **Strict**: A level that should be used when in production.

Note   • The bind DN is the string that the Cisco Cloud Network Controller uses to log in to the LDAP server. The Cisco Cloud Network Controller 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 Cisco Cloud Network Controller searches for the remote user account. This is where the password is validated. Filter is used to locate the attribute that the Cisco Cloud Network Controller requests to use for the *cisco-av-pair*. This contains the user authorization and assigned RBAC roles for use on the Cisco Cloud Network Controller. The Cisco Cloud Network Controller 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.

a) On the menu bar, choose **Administrative** > **Authentication**.

b) In the Work pane, click on **Login Domains** tab and then click on the **Actions** drop-down and select **Create Login Domain**.

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

Properties	Description
General	
Name	Enter the name of the Login Domain
Description	Enter the description of the Login Domain.
Settings	
Realm	Choose **LDAP** from the dropdown menu
Providers	To choose a Provider(s): 1. Click **Add Providers**. The **Select Providers** dialog appears. 2. Click to choose a provider(s) in the column on the left. 3. Click **Select**. You return to the **Create Login Domain** dialog box.

Properties	Description

Authentication Type	1. Select **Cisco AV Pairs**, if provider(s) was configured with **CiscoAVPair** as the **Attribute**.
	
	2. Select **LDAP Group Map Rules**, if provider(s) was configured with **memberOf** as the **Attribute**.
	
	   a. Click **Add LDAP Group Map Rule**. The dialog box appears.
	
	   b. Specify the map rule **Name**, **Description** (optional), and **Group DN**.
	
	   c. Click the + next to **Add Security Domain**. The dialog box appears.
	
	   d. Select the security domain using the **Select Security Domain** option.
	
	   e. Click the + to access the **Role** name and Role **Privilege** Type (**Read** or **Write**) fields. Click check mark.
	
	   f. If necessary, repeat the previous step to add more roles. Then click **Add**.
	
	   g. If you want to add more security domains, click the + next to **Add Security Domain**, then follow those steps again. Then click **Add**.

d) Click **Save** on Create Login Domain dialog box.

Configuring Cisco Cloud Network Controller for SAML Access

The following sections provide detailed information on configuring Cisco Cloud Network Controller 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 Cisco Cloud Network Controller for SAML Access

Note   SAML based Authentication is only for Cisco Cloud Network Controller 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 Cisco Cloud Network Controller 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 Cisco Cloud Network Controller, 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:

   - Choose the **Identity Provider** option (**ADFS**, **OKTA**, or **PING IDENTITY**).

   - 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 URL** 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.

   - Enter a value in the **GUI Redirect Banner Message (URL)** field.

   - Select the **Certificate Authority** if IdP is signed by a Private CA.

   - Enter a value in the **Timeout (sec)** field.

- Enter a value in the **Retries** field.

- Select the **Signature Algorithm Authentication User Requests** from the drop-down.

- Select checkbox to enable **Sign SAML Authentication Requests**, **Sign SAML Response Message**, **Sign Assertions in SAML Response**, **Encrypt SAML Assertions**.

g) Click **Save** to save the configuration.

Step 2 Create the login domain for SAML.

a) On the menu bar, choose **Administrative** > **Authentication**.

b) In the **Work** pane, click on the **Login Domains** tab and then click on the **Actions** drop-down and select **Create Login Domain**.

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

Properties	Description
General	
Name	Enter the name of the Login Domain
Description	Enter the description of the Login Domain.
Settings	
Realm	Choose **SAML** from the dropdown menu
Providers	To choose a Provider(s): 1. Click **Add Providers**. The **Select Providers** dialog appears. 2. Click to choose a provider(s) in the column on the left. 3. Click **Select**. You return to the **Create Login Domain** dialog box.

d) Click **Save** to save the configuration.

Setting Up a SAML Application in Okta

Refer to the section *Setting Up a SAML Application in Okta* of ***Cisco APIC Security Configuration Guide, Release 4.0(1)*** at
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html.

Setting Up a Relying Party Trust in AD FS

Refer to the section *Setting Up a Relying Party Trust in AD FS* in the **Cisco APIC Security Configuration Guide, Release 4.0(1)** at
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html.

Configuring HTTPS Access

The following sections describe how to configure HTTPS access.

About HTTPS Access

This article provides an example of how to configure a custom certificate for HTTPS access when using Cisco ACI.

For more information, see the section *HTTPS Access* in the **Cisco APIC Security Configuration Guide, Release 4.0(1)** at https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/Cisco-APIC-Security-Configuration-Guide-401.html.

Guidelines for Configuring Custom Certificates

- Wild card certificates (such as *.cisco.com, which is used across multiple devices) and its associated private key generated elsewhere are not supported on the Cisco Cloud Network Controller as there is no support to input the private key or password in the Cisco Cloud Network Controller. 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 Network Controller 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 Network Controller.

    - 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 Cloud Network Controller 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 Cisco Cloud Network Controller 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 (Cisco Cloud Network Controller). 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 Cisco Cloud Network Controller.

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 Cisco Cloud Network Controller.