



Managing Google Cloud Sites Using Nexus Dashboard Orchestrator

First Published: 2021-12-17

Last Modified: 2021-12-17

Americas Headquarters

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

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

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at: <http://www.cisco.com/go/softwareterms>. Cisco product warranty information is available at <http://www.cisco.com/go/warranty>. US Federal Communications Commission Notices are found here <http://www.cisco.com/c/en/us/products/us-fcc-notice.html>.

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

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

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

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

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

© 2021 Cisco Systems, Inc. All rights reserved.



CONTENTS

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

CHAPTER 2	Overview	3
	Google Cloud Overview	3
	Locating Important Google Cloud Project Information	3
	Understanding Google Cloud Deployments with Cloud APIC	4
	Inter-Site Connectivity Using BGP-EVPN	5
	External Network Connectivity	8
	Configuring Routing and Security Policies Separately	9
	Configuring Routing Policies	9
	Configuring Security Policies	10

CHAPTER 3	Configuring BGP-EVPN Intersite Connectivity	13
	Configuring Infra: Orchestrator General Settings	13
	Refreshing Cloud Site Connectivity Information	16
	Configuring Infra: Google Cloud Site Settings	17

CHAPTER 4	Configuring External Connectivity	21
	Configuring Google Cloud Site Connectivity Workflow	21
	Creating External VRFs in Infra Tenant	22
	Configuring Inter-site Connectivity Between Google Cloud Site and On-Premises Sites	23
	Adding External Devices	23
	Establishing Intersite Connectivity Between Google Cloud Site and On-Premises Sites	25
	Deploying Configuration to External Devices	27
	Configuring Intersite Connectivity Between Google Cloud Site and Other Cloud Sites	29

Deploying Infra Configuration	33
Creating an External EPG	34
Importing Google Cloud User Tenant	34
Creating a Tenant	35
Setting Up the Google Cloud Project for a User Tenant	36
Generating and Downloading Private Key Information from Google Cloud for an Unmanaged Tenant	38
Creating Google Cloud User Tenant	38
Setting the Necessary Permissions in Google Cloud for a Managed Tenant	41
Creating Cloud EPGs	42
Creating Schema, Template and VRFs for your Google Cloud Site	43
Configuring an Application Profile and EPG	44
Adding Cloud Endpoint Selector	44
Applying Contract Between External EPG and Cloud EPG	46
Configuring Route Leaking Between Cloud VRF and External VRF	47

CHAPTER 5

Configuring Internal Connectivity for Google Cloud Workloads	49
Internal Connectivity Workflow	49
Importing Google Cloud User Tenant	49
Creating a Tenant	50
Setting Up the Google Cloud Project for a User Tenant	50
Generating and Downloading Private Key Information from Google Cloud for an Unmanaged Tenant	52
Creating Google Cloud User Tenant	53
Setting the Necessary Permissions in Google Cloud for a Managed Tenant	56
Creating Schema, Template and VRFs for your Google Cloud Site	58
Creating Cloud EPGs	58
Applying contract between the cloud EPGs	59
Configuring Route Leaking between Two Cloud VRFs	60



CHAPTER 1

New and Changed Information

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

New and Changed Information

The following table provides an overview of the significant changes to the organization and features in this guide from the release the guide was first published to the current release. The table does not provide an exhaustive list of all changes made to the guide.

Table 1: Latest Updates

Release	New Feature or Update	Where Documented
August 28, 2022	Nexus Dashboard Orchestrator release 4.0(2) adds support for BGP-EVPN intersite connectivity for Google Cloud sites.	Inter-Site Connectivity Using BGP-EVPN , on page 5 and Configuring BGP-EVPN Intersite Connectivity , on page 13.
December 18, 2021	First release of this document.	--



CHAPTER 2

Overview

- [Google Cloud Overview, on page 3](#)
- [Inter-Site Connectivity Using BGP-EVPN, on page 5](#)
- [External Network Connectivity, on page 8](#)
- [Configuring Routing and Security Policies Separately, on page 9](#)

Google Cloud Overview

The following sections provide a brief overview of Google Cloud concepts as they relate to the Cisco Cloud APIC and Nexus Dashboard Orchestrator. For detailed information about Cloud APIC deployment and configuration, see the [Cloud APIC documentation](#).

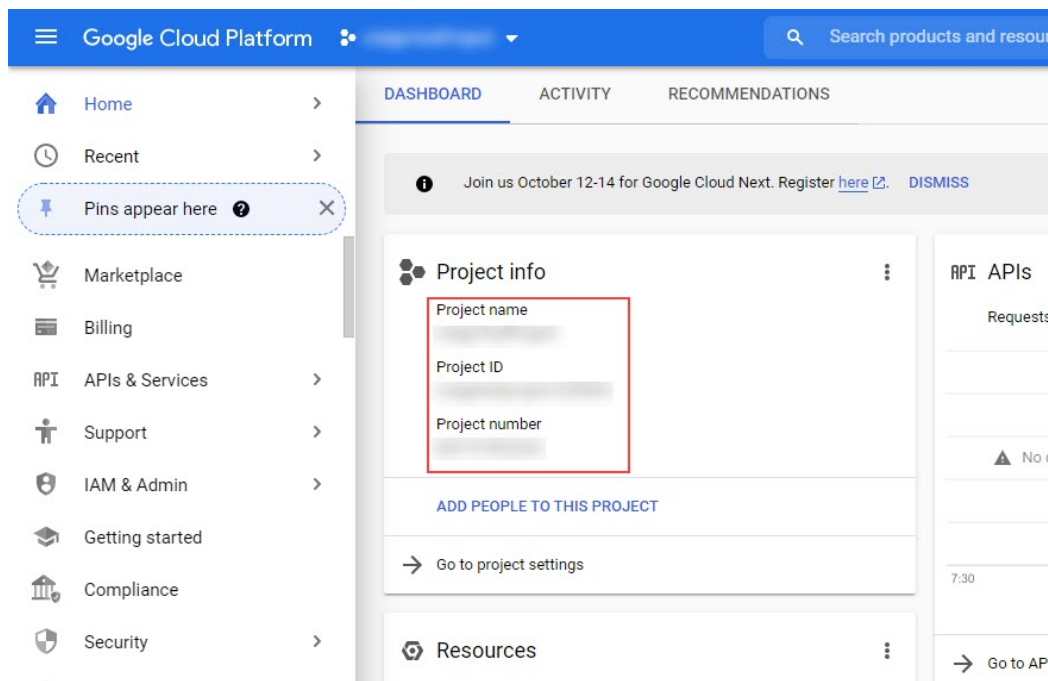
Locating Important Google Cloud Project Information

The following information is required if you plan to create new tenants in your Google Cloud sites. If you plan to import existing tenants only, you can skip this section.

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 Google Cloud 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 Cloud APIC

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 Cloud APIC perspective.

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

- A Cloud APIC tenant cannot span multiple Google Cloud Projects
- There cannot be more than one Cloud APIC tenant in a Google Cloud Project

With Cloud APIC, 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 Cloud APIC 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 Cloud APIC, other tenants).

User Tenants With Managed Credentials

This type of user tenant has the following characteristics:

- This tenant account is managed by the Cisco Cloud APIC.

- You will first choose **Managed Identity** in the Nexus Dashboard Orchestrator GUI as part of the tenant configuration process for this type of user tenant.
- After you have configured the necessary parameters in the Nexus Dashboard Orchestrator, you must then set the necessary roles for this tenant in Google Cloud. Add the service account created by the Cloud APIC 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

User Tenants With Unmanaged Credentials

This type of user tenant has the following characteristics:

- This tenant account is not managed by the Cisco Cloud APIC.
- Before configuring the necessary parameters in the Cisco Cloud APIC 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 Nexus Dashboard Orchestrator 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 Nexus Dashboard Orchestrator, you will provide the following information from the downloaded JSON file:
 - Key ID
 - RSA Private Key
 - Client ID
 - Email

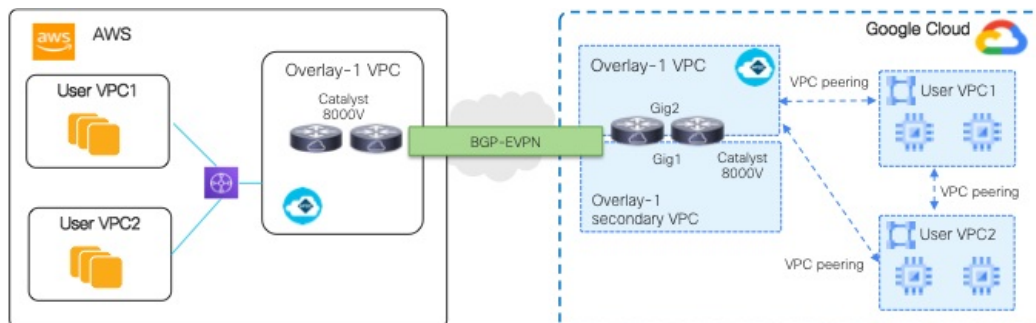
Inter-Site Connectivity Using BGP-EVPN

Beginning with Cloud Network Controller release 25.0(5), 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.



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.

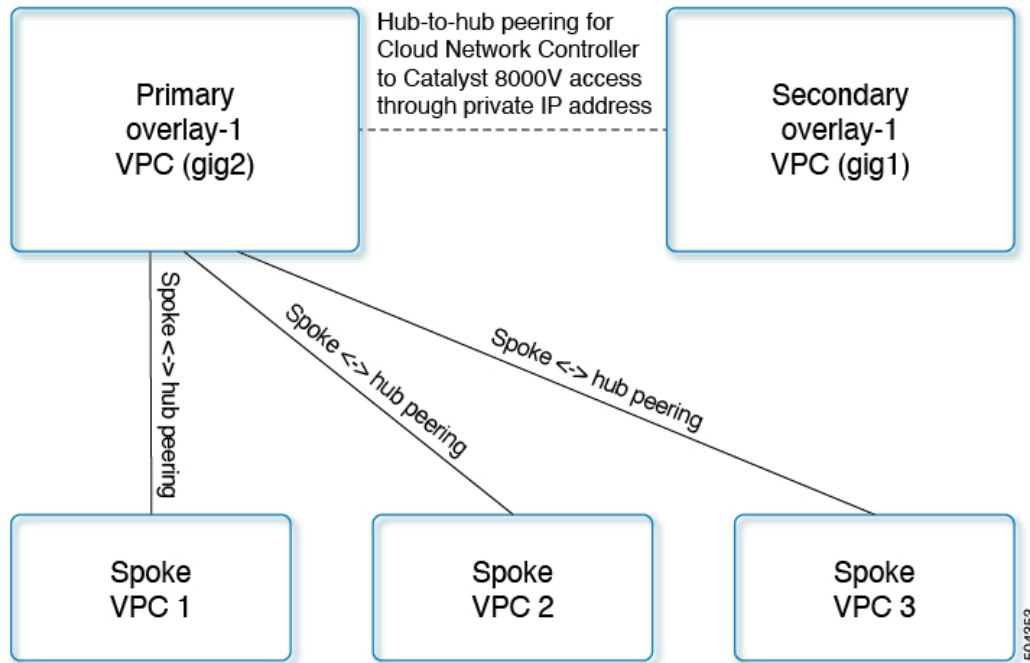
VPC Peering

In order to have communication 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:

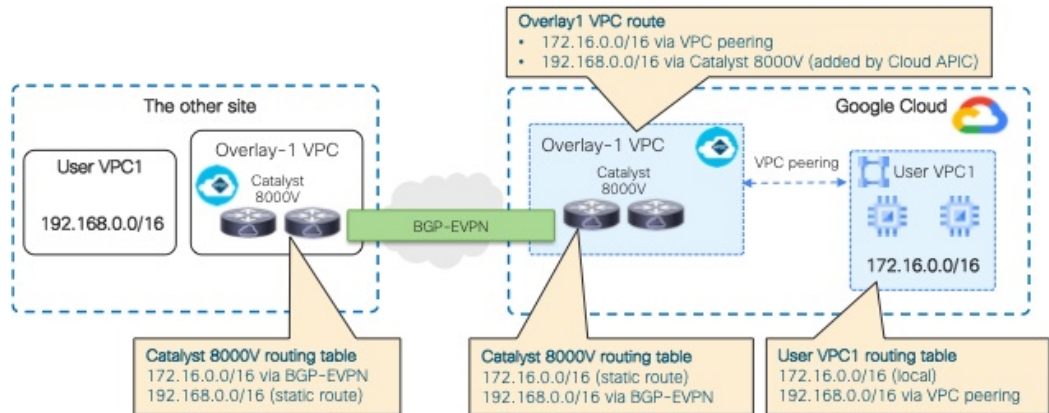
- Spoke-to-spoke VPC peering: This is used for spoke-to-spoke intra-site communication.
- Hub-to-spoke VPC peering: This is used for inter-site communication that goes through the Cisco Catalyst 8000V routers using BGP-EVPN.
- Hub-to-hub VPC peering: This is used for communication between the Cisco Cloud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

External Network Connectivity

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, including a CSR.

The following sections provide more information on the components that allow for the new external network connectivity provided in Cloud APIC release 25.0(2) and later.

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 Nexus Dashboard Orchestrator.

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 APIC 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 BGP-IPv4 connectivity is supported for this type of connectivity using cloud native routers, where BGP-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 APIC 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 APIC has a restriction where only one cloud router and one cloud VPN gateway is supported per region.

VPN Communication

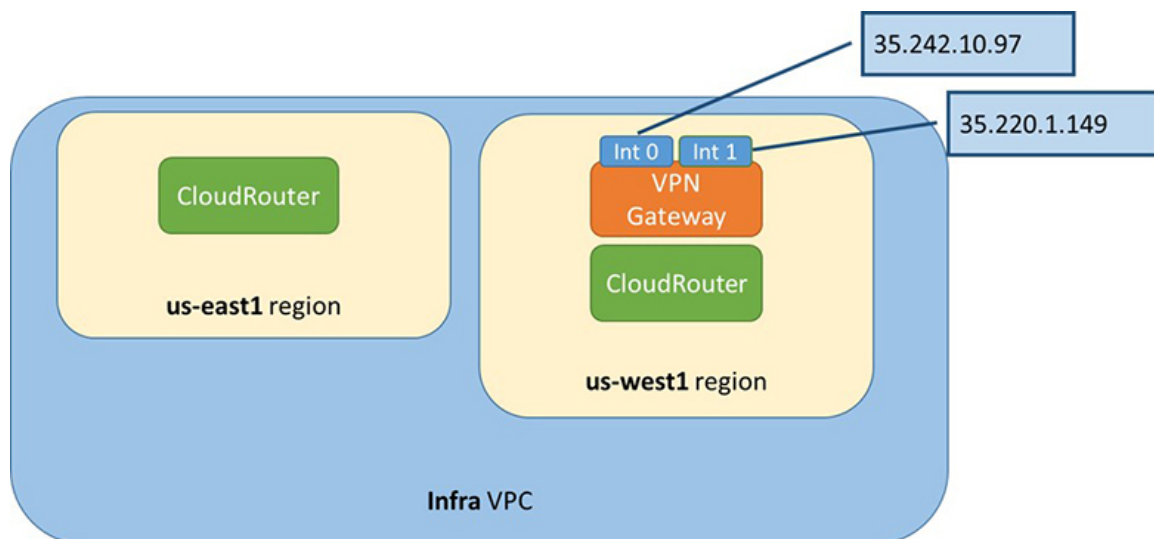
When configuring Cisco Cloud APIC with Google Cloud, the infra VPC is used to host the Cisco Cloud APIC and to host the VPN connections to external devices and sites. However, the infra VPC is not used as a transit to implement spoke-to-spoke communication. Instead, when configuring Cisco Cloud APIC with Google Cloud, spoke-to-spoke communication is done through 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 APIC only supports VPN gateways with two interfaces because two interfaces are required to achieve high availability.



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 VRF route leaking.

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 (t_1 and t_2) and two corresponding VRFs (v_1 and v_2). For every route leak entry $t_1:v_1$ under the VRF $t_2:v_2$, there must be a corresponding route leak entry $t_2:v_2$ under the VRF $t_1:v_1$.
- 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 APIC 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 APIC Implements Firewall Rules With Google Cloud

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

- **Global resources:** VPCs and firewalls in Google Cloud are global resources, so Cisco Cloud APIC 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 APIC configures a pair of low-priority deny-all ingress and egress rules when the VPC is created. Afterwards, Cisco Cloud APIC 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 APIC, 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 APIC 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 APIC 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 APIC recognizes the labels assigned to the VMs in Google Cloud.

Cisco Cloud APIC 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 APIC 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.

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

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

Endpoint	Application Profile	EPG	Primary IP	Labels	Cloud APIC-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 APIC 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 APIC 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 APIC 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 APIC:** 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.



CHAPTER 3

Configuring BGP-EVPN Intersite Connectivity

- [Configuring Infra: Orchestrator General Settings, on page 13](#)
- [Refreshing Cloud Site Connectivity Information, on page 16](#)
- [Configuring Infra: Google Cloud Site Settings, on page 17](#)

Configuring Infra: Orchestrator General Settings

This section describes how to configure general Infra settings for all the sites.



Note Some of the following settings apply to all sites, while others are required for specific type of sites (for example, Cloud Network Controller sites). Ensure that you complete all the required configurations in infra general settings before proceeding to the site-local settings specific to each site.

-
- Step 1** Log in to the Cisco Nexus Dashboard Orchestrator GUI.
- Step 2** In the left navigation menu, select **Infrastructure > Site Connectivity**.
- Step 3** In the main pane, click **Configure**.
- Step 4** In the left sidebar, select **General Settings**.
- Step 5** Provide **Control Plane Configuration**.
- a) Select the **Control Plane Configuration** tab.
 - b) Choose **BGP Peering Type**.
 - **full-mesh**—All border gateway switches in each site will establish peer connectivity with remote sites' border gateway switches.
In **full-mesh** configuration, Nexus Dashboard Orchestrator uses the spine switches for ACI managed fabrics and border gateways for NDFC managed fabrics.
 - **route-reflector**—The route-reflector option allows you to specify one or more control-plane nodes to which each site establishes MP-BGP EVPN sessions. The use of route-reflector nodes avoids creating MP-BGP EVPN full mesh adjacencies between all the sites managed by NDO.
For ACI fabrics, the **route-reflector** option is effective only for fabrics that are part of the same BGP ASN.
 - c) In the **Keepalive Interval (Seconds)** field, enter the keep alive interval seconds.

We recommend keeping the default value.

- d) In the **Hold Interval (Seconds)** field, enter the hold interval seconds.

We recommend keeping the default value.

- e) In the **Stale Interval (Seconds)** field, enter stale interval seconds.

We recommend keeping the default value.

- f) Choose whether you want to turn on the **Graceful Helper** option.

- g) Provide the **Maximum AS Limit**.

We recommend keeping the default value.

- h) Provide the **BGP TTL Between Peers**.

We recommend keeping the default value.

- i) Provide the **OSPF Area ID**.

If you do not have any Cloud Network Controller sites, this field will not be present in the UI.

This is OSPF area ID used by cloud sites for on-premises IPN peering.

- j) (Optional) Enable **IANA Assigned Port** for CloudSec encryption.

By default, CloudSec uses a proprietary UDP port. This option allows you to configure CloudSec to use the official IANA-reserved port 8017 for CloudSec encryption between sites.

Note The IANA-reserved port is supported for Cisco APIC sites running release 5.2(4) or later.

To change this setting, CloudSec must be disabled on all sites. If you want to enable IANA reserved port, but already have CloudSec encryption enabled for one or more of your sites, disable CloudSec for all sites, enable **IANA Reserve UDP Port** option, then re-enable CloudSec for the required sites.

For detailed information and steps for configuring CloudSec, see the "CloudSec Encryption" chapter of the [Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics](#).

Step 6 Provide the **IPN Devices** information.

If you do not plan to configure inter-site connectivity between on-premises and cloud sites, you can skip this step.

When you configure inter-site underlay connectivity between on-premises and cloud sites as described in later sections, you will need to select an on-premises IPN device which will establish connectivity to the cloud CSRs. These IPN devices must first be defined here before they are available in the on-premises site configuration screen.

- Select the **On Premises IPsec Devices** tab.
- Click **+Add On-Premises IPsec Device**.
- Choose whether the device is **Unmanaged** or **Managed** and provide the device information.

This defines whether or not the device is directly managed by NDFC:

- For **Unmanaged** IPN devices, simply provide the **Name** and the **IP Address** of the device.

The IP address you provide will be used as the tunnel peer address from the cloud CSRs, not the IPN device's management IP address.

- For **Managed** IPN devices, choose the NDFC **Site** that contains the device and then the **Device** from that site.

Then choose the **Interface** on the device that is facing the Internet and provide the **Next Hop** IP address, which is the IP address of the gateway that is connecting to the Internet.

- d) Click the check mark icon to save the device information.
- e) Repeat this step for any additional IPN devices you want to add.

Step 7 Provide the **External Devices** information.

If you do not have any Cloud Network Controller sites, this tab will not be present in the UI.

If you do not have any Cloud Network Controller sites in your Multi-Site domain or you do not plan to configure connectivity between cloud sites and branch routers or other external devices, you can skip this step.

The following steps describe how to provide information about any branch routers or external devices to which you want to configure connectivity from your cloud sites.

- a) Select the **External Devices** tab.

This tab will only be available if you have at least one cloud site in your Multi-Site domain.

- b) Click **Add External Device**.

The **Add External Device** dialogue will open.

- c) Provide the **Name**, **IP Address**, and **BGP Autonomous System Number** for the device.

The IP address you provide will be used as the tunnel peer address from the Cloud Network Controller's CSRs, not the device's management IP address. The connectivity will be established over public Internet using IPsec.

- d) Click the check mark icon to save the device information.
- e) Repeat this step for any additional IPN devices you want to add.

After you have added all the external devices, ensure to complete the next step to provide the IPsec tunnel subnet pools from with the internal IP addresses will be allocated for these tunnels.

Step 8 Provide the **IPsec Tunnel Subnet Pools** information.

If you do not have any Cloud Network Controller sites, this tab will not be present in the UI.

There are two types of subnet pools that you can provide here:

- **External Subnet Pool**—used for connectivity between cloud site CSRs and other sites (cloud or on-premises).

These are large global subnet pools that are managed by Nexus Dashboard Orchestrator. The Orchestrator, creates smaller subnets from these pools and allocates them to sites to be used for inter-site IPsec tunnels and external connectivity IPsec tunnels.

You must provide at least one external subnet pool if you want to enable external connectivity from one or more of your cloud sites.

- **Site-Specific Subnet Pool**—used for connectivity between cloud site CSRs and external devices.

These subnets can be defined when the external connectivity IPsec tunnels must be in a specific range. For example, where a specific subnet is already being used to allocate IP addresses to the external router and you want to continue using those subnets for IPsec tunnels for NDO and cloud sites. These subnets are not managed by the Orchestrator and each subnet is assigned to a site in its entirety to be used locally for external connectivity IPsec tunnels.

If you do not provide any named subnet pools but still configure connectivity between cloud site's CSRs and external devices, the external subnet pool will be used for IP allocation. .

Note The minimum mask length for both subnet pools is /24.

To add one or more **External Subnet Pools**:

- a) Select the **IPSec Tunnel Subnet Pools** tab.
- b) In the **External Subnet Pool** area, click **+Add IP Address** to add one or more external subnet pools.

This subnet will be used to address the IPsec tunnel interfaces and loopbacks of the Cloud Routers used for on-premises connectivity, which you previously configured in the Cloud Network Controller for inter-site connectivity in earlier Nexus Dashboard Orchestrator releases.

The subnets must not overlap with other on-premises TEP pools, should not begin with `0.x.x.x` or `0.0.x.x`, and should have a network mask between `/16` and `/24`, for example `30.29.0.0/16`.

- c) Click the check mark icon to save the subnet information.
- d) Repeat these substeps for any additional subnet pools you want to add.

To add one or more **Site-Specific Subnet Pools**:

- a) Select the **IPSec Tunnel Subnet Pools** tab.
- b) In the **Site-Specific Subnet Pools** area, click **+Add IP Address** to add one or more external subnet pools.

The **Add Named Subnet Pool** dialogue will open.

- c) Provide the subnet **Name**.

You will be able to use the subnet pool's name to choose the pool from which to allocate the IP addresses later on.

- d) Click **+Add IP Address** to add one or more subnet pools.

The subnets must have a network mask between `/16` and `/24` and not begin with `0.x.x.x` or `0.0.x.x`, for example `30.29.0.0/16`.

- e) Click the check mark icon to save the subnet information.

Repeat the steps if you want to add multiple subnets to the same named subnet pool.

- f) Click **Save** to save the named subnet pool.
- g) Repeat these substeps for any additional named subnet pools you want to add.

What to do next

After you have configured general infra settings, you must still provide additional information for site-specific configurations based on the type of sites (ACI, Cloud Network Controller, or NDFC) you are managing. Follow the instructions described in the following sections to provide site-specific infra configurations.

Refreshing Cloud Site Connectivity Information

Any infrastructure changes, such as CSR and Region addition or removal, require a Multi-Site fabric connectivity site refresh. This section describes how to pull up-to-date connectivity information directly from each site's controller.

-
- Step 1** Log in to the Cisco Nexus Dashboard Orchestrator GUI.
 - Step 2** In the left navigation menu, select **Infrastructure > Site Connectivity**.
 - Step 3** In the top right of the main pane, click **Configure**.
 - Step 4** In the left pane, under **Sites**, select a specific site.

- Step 5** In the main window, click the **Refresh** button to discover any new or changed CSRs and regions.
- Step 6** Finally, click **Yes** to confirm and load the connectivity information.
This will discover any new or removed CSRs and regions.
- Step 7** Click **Deploy** to propagate the cloud site changes to other sites that have connectivity to it.
After you refresh a cloud site's connectivity and CSRs or regions are added or removed, you need to deploy infra configuration so other sites that have underlay connectivity to that cloud site get updated configuration.
-

Configuring Infra: Google Cloud Site Settings

This section describes how to configure site-specific Infra settings for Cloud Network Controller sites.

- Step 1** Log in to the Cisco Nexus Dashboard Orchestrator GUI.
- Step 2** In the left navigation menu, select **Infrastructure > Site Connectivity**.
- Step 3** In the top right of the main pane, click **Configure**.
- Step 4** In the left pane, under **Sites**, select a specific cloud site.
- Step 5** Provide the general **Inter-Site Connectivity** information.
- In the right **<Site> Settings** pane, select the **Inter-Site Connectivity** tab.
 - Enable the **Multi-Site** knob.
This defines whether the overlay connectivity is established between this site and other sites.
Note that the overlay configuration will not be pushed to sites which do not have the underlay intersite connectivity established as described in the next step.
- Step 6** Provide site-specific **Inter-Site Connectivity** information.
- If using the BGP-EVPN protocol for site connectivity, enable **Contract Based Routing** option.
 - In the right properties sidebar for the cloud site, click **Add Site**.
The **Add Site** window opens.
 - Under **Connected to Site**, click **Select a Site** and select the site (for example, `site2`) to which you want to establish connectivity from the site you are configuring (for example, `site1`).
Once you select the remote site, the **Add Site** window will update to reflect both directions of connectivity: **Site1 > Site2** and **Site2 > Site1**.
 - In the **Site1 > Site2** area, from the **Connection Type** dropdown, choose the type of connection between the sites.
The following options are available:
 - Public Internet**—connectivity between the two sites is established via the Internet.
This type is supported between any two cloud sites or between a cloud site and an on-premises site.
 - Private Connection**—connectivity is established using a private connection between the two sites.
This type is supported between a cloud site and an on-premises site.

- `Cloud Backbone`—connectivity is established using cloud backbone.

This type is supported between two cloud sites of the same type, such as Azure-to-Azure, AWS-to-AWS, or GCP-to-GCP.

If you have multiple types of sites (on-premises, AWS, Azure, and GCP), different pairs of site can use different connection type.

- e) Choose the **Protocol** that you want to use for connectivity between these two sites.

For this use case, we will use **BGP-EVPN**. You can optionally enable **IPSec** and choose which version of the Internet Key Exchange (IKE) protocol to use: IKEv1 (`Version 1`) or IKEv2 (`Version 1`) depending on your configuration.

- For `Public Internet` connectivity, IPsec is always enabled.
- For `Cloud Backbone` connectivity, IPsec is always disabled.
- For `Private Connection`, you can choose to enable or disable IPsec.

If using **BGP-IPv4** connectivity instead, you must provide an external VRF which will be used for route leaking configuration from the cloud site you are configuring.

After **Site1 > Site2** connectivity information is provided, the **Site2 > Site1** area will reflect the connectivity information in the opposite direction.

- f) Click **Save** to save the inter-site connectivity configuration.

When you save connectivity information from `Site1` to `Site2`, the reverse connectivity is automatically created from `Site2` to `Site1`, which you can see by selecting the other site and checking the **Inter-site Connectivity** information in the right sidebar.

- g) Repeat this step to add inter-site connectivity for other sites.

When you establish underlay connectivity from `Site1` to `Site2`, the reverse connectivity is done automatically for you.

However, if you also want to establish inter-site connectivity from `Site1` to `Site3`, you must repeat this step for that site as well.

Step 7 Provide **External Connectivity** information.

If you do not plan to configure connectivity to external sites or devices that are not managed by NDO, you can skip this step.

Detailed description of an external connectivity use case is available in the [Configuring External Connectivity from Cloud CSRs Using Nexus Dashboard Orchestrator](#) document.

- In the right `<Site> Settings` pane, select the **External Connectivity** tab.
- Click **Add External Connection**.

The **Add External Connectivity** dialog will open.

- From the **VRF** dropdown, select the VRF you want to use for external connectivity.

This is the VRF which will be used to leak the cloud routes. The **Regions** section will display the cloud regions that contain the CSRs to which this configuration be applied.

- From the **Name** dropdown in the **External Devices** section, select the external device.

This is the external device you added in the **General Settings > External Devices** list during general infra configuration and must already be defined as described in [Configuring Infra: Orchestrator General Settings, on page 13](#).

- e) From the **Tunnel IKE Version** dropdown, pick the IKE version that will be used to establish the IPsec tunnel between the cloud site's CSRs and the external device.
- f) (Optional) From the **Tunnel Subnet Pool** dropdown, choose one of the named subnet pools.

Named subnet pools are used to allocate IP addresses for IPsec tunnels between cloud site CSRs and external devices. If you do not provide any **named** subnet pools here, the **external** subnet pool will be used for IP allocation.

Providing a dedicated subnet pool for external device connectivity is useful for cases where a specific subnet is already being used to allocate IP addresses to the external router and you want to continue to use those subnets for IPsec tunnels for NDO and cloud sites.

If you want to provide a specific subnet pool for this connectivity, it must already be created as described in [Configuring Infra: Orchestrator General Settings, on page 13](#).

- g) (Optional) In the **Pre-Shared Key** field, provide the custom keys you want to use to establish the tunnel.
- h) If necessary, repeat the previous substeps for any additional external devices you want to add for the same external connection (same VRF).
- i) If necessary, repeat this step for any additional external connections (different VRFs).

Note that there's a one-to-one relationship for tunnel endpoints between CSRs and external devices, so while you can create additional external connectivity using different VRFs, you cannot create additional connectivity to the same external devices.



CHAPTER 4

Configuring External Connectivity

- [Configuring Google Cloud Site Connectivity Workflow, on page 21](#)
- [Creating External VRFs in Infra Tenant, on page 22](#)
- [Configuring Inter-site Connectivity Between Google Cloud Site and On-Premises Sites, on page 23](#)
- [Configuring Intersite Connectivity Between Google Cloud Site and Other Cloud Sites, on page 29](#)
- [Deploying Infra Configuration, on page 33](#)
- [Creating an External EPG, on page 34](#)
- [Importing Google Cloud User Tenant, on page 34](#)
- [Creating a Tenant, on page 35](#)
- [Creating Cloud EPGs, on page 42](#)
- [Applying Contract Between External EPG and Cloud EPG , on page 46](#)
- [Configuring Route Leaking Between Cloud VRF and External VRF, on page 47](#)

Configuring Google Cloud Site Connectivity Workflow

The following sections describe how to configure Google Cloud sites infra, intersite connectivity, and a simple deployment use case. The workflow includes:

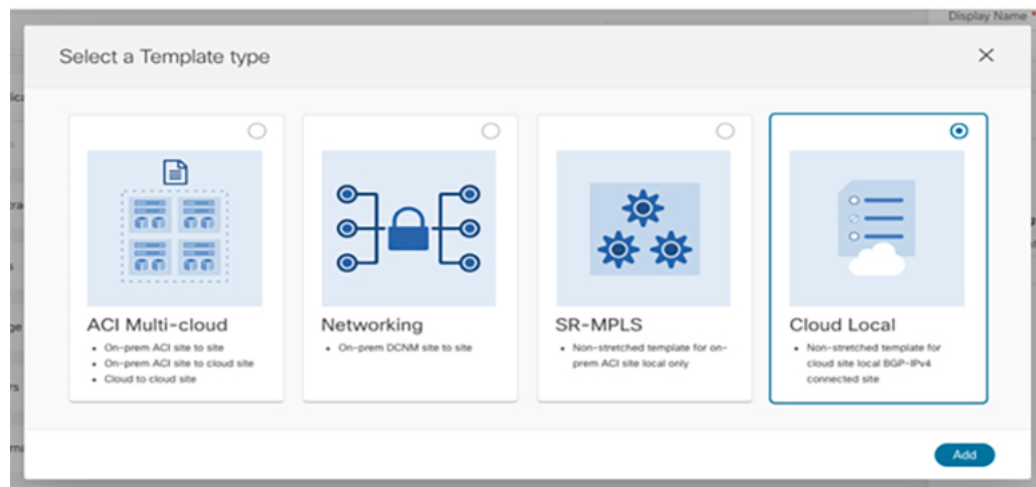
- Configuring general infra settings, such as adding the on-premises IPN devices as external devices in Nexus Dashboard Orchestrator and establishing external connectivity from google cloud site to those devices
- Configuring and deploying external VRFs in the Google Cloud site's Infra tenant
- Configuring intersite connectivity from your Google Cloud site to an on-premises site and manually configuring any on-premises sites connectivity to the Google Cloud site
- Configuring intersite connectivity between your Google Cloud site and other cloud sites like AWS/Azure
- Configuring route leaking in the external VRFs to enable routing between sites
- Creating or importing a user tenant and EPGs and applying contracts to enable communication between sites

Creating External VRFs in Infra Tenant

You can create a single schema where you will define the external VRFs for all cloud sites in your multi-Site domain. However, since you may want to deploy different VRFs for each cloud site, you must create separate templates for each cloud site you have because templates cannot be shared across different cloud sites.

Following sections will introduce you to the new type of template and will walk you through the process of adding external VRFs. In the schema, create new template for Google Cloud site, use cloud local template, assign the template to the Google Cloud site.

You will have an option to a new type of template which is called Cloud Local.



This type of template cannot be stretched to multiple sites. It supports and allow all types of cloud sites. However, no more than one site can be attached to this template. There is another restriction regarding this template which is some of the objects should be only from within the tenant, Like VRF etc.

The following section describes how to create an external VRF which will be used to establish connectivity to an external devices' subnets. You can follow the provided steps to provision an external VRF to a cloud site.

-
- Step 1** Log in to the Cisco Nexus Dashboard Orchestrator GUI.
- Step 2** In the **Main menu**, select **Application Management > Schemas**.
- Step 3** Create a new schema and templates or select an existing schema where you will deploy the templates associated to the Infra tenant containing the external VRFs definition.

You can create a separate schema specifically for this use case, where you will define all templates associated to the Infra tenant and containing the external VRFs providing the connectivity to the external devices.

When creating the external VRF templates:

- You must use separate templates for different types of cloud sites (AWS, Azure or Google Cloud).
- You must choose the **Cloud Local** template type.
- You must map the template to the `infra` tenant or the VRFs cannot be used for external connectivity.
- You can use the same VRF name in both templates. We will use `extVrf1` for the examples in this document.

Step 4 In the main pane, select **+Create Object > VRF**.

Step 5 Provide the **Display Name** for the VRF.

You can leave all other options at default values.

Note In the VRF's site local properties, do not attach this VRF to any regions. Any VRF that is created in the Infra tenant and is not attached to any region is treated as an external VRF and can be used for this use case.

Step 6 Assign the template that contains your external VRF to one or more cloud sites from which you will establish external connectivity.

Remember that you must assign the template only to one type of cloud sites (AWS, Azure or Google Cloud).

Step 7 Deploy the templates to create the external VRF in the cloud sites.

Configuring Inter-site Connectivity Between Google Cloud Site and On-Premises Sites

The following sections describe how to configure connectivity between your Google Cloud site and an on-premises site. If you want to configure connectivity between two cloud sites, see [Configuring Intersite Connectivity Between Google Cloud Site and Other Cloud Sites](#), on page 29.

Adding External Devices

If you do not plan to establish intersite connectivity between you Google Cloud site and an on-premises site, you may skip this section. This section describes how to provide information about your external devices to your Nexus Dashboard Orchestrator in the Orchestrator's **Site Connectivity** page.



Note The following steps focus on the configurations required for this specific use case. Detailed information about all infra configuration settings is available in [Cisco Nexus Dashboard Orchestrator Configuration Guides](#).

Step 1 Log in to the Cisco Nexus Dashboard Orchestrator GUI.

Step 2 In the left navigation menu, select **Infrastructure > Site Connectivity**.

Step 3 In the main pane, click **Configure**.

Step 4 In the left sidebar, select **General Settings**.

Step 5 Provide the **External Devices** information.

This step describes how to provide information about any external devices to which you want to configure connectivity from your cloud sites, for more information regarding this process, see [Configure General Infra Settings](#).

- a) Select the **External Devices** tab.
- b) Click **Add External Device**.

The **Add External Device** dialogue will open.

- c) Provide the **Name**, **IP Address**, and **BGP Autonomous System Number** for the device.

The IP address you provide will be used as peer address for the Cloud APIC CSRs or Google Cloud native router's VPN gateway, not the device's management IP address. The connectivity will be established over public Internet using IPsec.

- d) Click **Save**, to save the device information.
e) Repeat this step for any additional external devices you want to add.

Step 6 Enter the necessary information in the **IPSec Tunnel Subnet Pools** area.

By default, a subnet pool of `169.254.0.0/16` is populated to create the IPsec tunnels between the Google cloud and other cloud sites (AWS/Azure). 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.

The following subnets are reserved and cannot be used for any tunnels:

- `169.254.0.0/30`
- `169.254.1.0/30`
- `169.254.2.0/30`
- `169.254.3.0/30`
- `169.254.4.0/30`
- `169.254.5.0/30`
- `169.254.112.0/24`
- `169.254.113.0/24`
- `169.254.114.0/24`
- `169.254.169.252/30`

There are two types of subnet pools that you can provide here:

- **External Subnet Pool**—used for connectivity between cloud site router and other sites (cloud or on-premises).

These are large global subnet pools that are managed by Nexus Dashboard Orchestrator. The Orchestrator, creates smaller subnets from these pools and allocates them to sites to be used for inter-site IPsec tunnels and external connectivity IPsec tunnels.

For AWS/Azure, you will have to provide at least one external subnet pool if you want to enable external connectivity from one or more of your cloud sites. However, For Google Cloud, you can leave the poolname blank (unselected) when configuring external device connectivity. In this case Nexus Dashboard Orchestrator will allocate a `/24` `169.254.0.0/16` from the subnet (from the top of the range, ie. it will be `169.254.255.0/24` etc).

- **Site-Specific Subnet Pool**—used for connectivity between cloud site router and external devices.

These subnets can be defined when the external connectivity IPsec tunnels must be in a specific range. For example, where a specific subnet is already being used to allocate IP addresses to the external router and you want to continue using those subnets for IPsec tunnels for Nexus Dashboard Orchestrator and cloud sites. These subnets are not managed by the Orchestrator and each subnet is assigned to a site in its entirety to be used locally for external connectivity IPsec tunnels.

You will assign a name to a site-specific subnet pool, such as `169.254.0.0/24`, which you can then use when configuring the external devices.

If you do not provide any named subnet pools but still configure connectivity between cloud site router and external devices, the external subnet pool will be used for IP allocation to the IPsec tunnels established between cloud site router and external devices.

Note The minimum mask length for both subnet pools is `/24`.

To add one or more **External Subnet Pools**:

- a) Select the **IPsec Tunnel Subnet Pools** tab.
- b) In the **External Subnet Pool** area, click **+Add IP Address** to add one or more external subnet pools.

This subnet will be used to address the IPsec tunnel interfaces and loopbacks of the Cloud Routers used for on-premises connectivity, which you previously configured in the Cloud APIC for inter-site connectivity in earlier Nexus Dashboard Orchestrator releases.

The subnets must not overlap with other on-premises TEP pools, should not begin with `0.x.x.x` or `0.0.x.x`, and should have a network mask between `/16` and `/24`, for example `10.12.0.0/16`.

- c) Click the check mark icon to save the subnet information.
- d) Repeat these substeps for any additional subnet pools you want to add.

To add one or more **Site-Specific Subnet Pool**:

- a) Select the **IPsec Tunnel Subnet Pools** tab.
- b) In the **Named Subnet Pool** area, click **+Add IP Address** to add one or more external subnet pools.

The **Add Named Subnet Pool** dialogue will open.

- c) Provide the subnet **Name**.

You will be able to use the subnet pool's name to choose the pool from which to allocate the IP addresses later on, for example `extSubPool1`.

- d) Click **+Add IP Address** to add one or more subnet pools.

The subnets must have a network mask between `/16` and `/24` and not begin with `0.x.x.x` or `0.0.x.x`, for example `10.181.0.0/16`.

- e) Click the check mark icon to save the subnet information.
Repeat the steps if you want to add multiple subnets to the same named subnet pool.
- f) Click **Save** to save the named subnet pool.
- g) Repeat these substeps for any additional named subnet pools you want to add.

Establishing Intersite Connectivity Between Google Cloud Site and On-Premises Sites

Before you begin, you must have:

- Created and deployed the external VRFs in your Google Cloud site, as described in [Creating External VRFs in Infra Tenant, on page 22](#).

- Added one or more external devices, as described in [Adding External Devices, on page 23](#).



Note Before configuring the external connectivity, you can refresh across all the sites in Fabric Connectivity Infra page and deploy to make sure that all CSRs for AWS/Azure and cloud routers for Google Cloud sites are reflected correctly in Nexus Dashboard Orchestrator.

Before you begin

This section describes how to configure site-specific Infra settings for Cloud APIC sites. Before starting make sure that you have:

- Created and deployed the external VRFs in your Google Cloud site, as described in [Creating External VRFs in Infra Tenant, on page 22](#).
- Added one or more external devices, as described in [Adding External Devices, on page 23](#).



Note Before configuring the external connectivity, you can refresh across all the sites in Fabric Connectivity Infra page and deploy to make sure that all CSRs for AWS/Azure and cloud routers for Google Cloud sites are reflected correctly in Nexus Dashboard Orchestrator.

Step 1 In the left pane of the **Fabric Connectivity Infra** page, under **Sites**, select a specific cloud site.

This is the site from which you want to establish connectivity to an external device.

Step 2 Provide **External Connectivity** information.

You must complete this step to provide connectivity information to the external devices as part of this use case configuration.

- In the right **<Site> Settings** pane, select the **External Connectivity** tab.
- Click **Add External Connection**.

The **Add External Connectivity** dialog will open.

- From the **VRF** dropdown, select the VRF you want to use for external connectivity.

This is the VRF (`extVrf1`) which will be used to leak the cloud routes and which you already created.

- Click **+Add External Device**.
- From the **Name** dropdown in the **External Devices** section, select the external device.

This is the external device which you added in the **General Settings > External Devices** list during general infra configuration and must already be defined.

- For the **Tunnel IKE Version**, IKE-V2 will be selected. As of this release, only IKE-V2 is supported.
- (Optional) From the **Tunnel Subnet Pool** dropdown, choose one of the site-specific subnet pools.

Site-specific subnet pools are used to allocate IP addresses for IPsec tunnels between cloud site router and external devices. If you do not provide any **Site-Specific Subnet Pool** subnet pools here, the **External Subnet Pool** subnet pool will be used for IP allocation.

Providing a dedicated subnet pool for external device connectivity is useful for cases where a specific subnet is already being used to allocate IP addresses to the external router and you want to continue to use those subnets for IPsec tunnels for Nexus Dashboard Orchestrator and cloud sites.

- h) (Optional) In the **Pre-Shared Key** field, provide the custom keys you want to use to establish the tunnel.
If you do not provide a pre-shared key, Cloud APIC will generate one automatically on the cloud site router.
- i) If necessary, repeat the previous substeps for any additional external devices you want to add for the same external connection (same external VRF).
- j) If necessary, repeat this step for any additional external connections (different external VRFs).

Note that there's a one-to-one relationship for tunnel endpoints between cloud site router and external devices, so while you can create additional external connectivity using different external VRFs, you cannot create additional connectivity to the same external devices.

Deploying Configuration to External Devices

While the previous section described how to deploy infra configuration to the cloud sites' Cloud APICs to enable connectivity from the cloud sites to the external devices, this section describes how to enable connectivity from the external device to the cloud sites.

Step 1 Gather the necessary information that you will need to enable connectivity from the external device.

You can get the required configuration details using either the **Deploy & Download External Device Config files** or the **Download External Device Config files** option in Nexus Dashboard Orchestrator as part of the procedure.

When you download the configuration files:

- The number of files will match the number of sites that have external connectivity.
- The file name affixes will match the site IDs.

For example, `<...>-2.config` indicates the file is for a site with Site ID 2. The site ID is listed in each site's **Site Connectivity** page in the Nexus Dashboard Orchestrator GUI.

Step 2 Log into the external device.

Step 3 Configure the tunnels and BGP from the external device to the cloud router.

When configuring external devices:

- Depending on the specific requirements, the external subnets may or may not be in the same VRF with the tunnel interfaces

If the external subnets are in different VRFs then proper route leaking must be configured on the external device

Note Note that the configuration downloaded from Nexus Dashboard Orchestrator only allows to establish IPsec and BGP connectivity. It does not provide any information on the route-leaking configuration within the external device itself.

- Once the external subnets are advertised to the cloud site router, Nexus Dashboard Orchestrator provisions the route leaking configuration to select the subnets to be imported into the user tenant VRF.
- The following examples assume BGP configuration is done in the external VRF (`extVrf1`) and the external subnets as well as tunnel interfaces on the external device are part of the same VRF.

The following example shows how to configure a single IPsec tunnel (`Tunnel100`) from an external device (in this case ASR1K) to a CSR:

Example:

```
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-ifc-7
  peer peer-ikev2-keyring
  address 35.220.81.45
  pre-shared-key 163988519666274287497025544399329641924
!
crypto ikev2 profile ikev-profile-ifc-7
  match address local interface GigabitEthernet1
  match identity remote address 35.220.81.45 255.255.255.255
  identity local address 20.92.217.94
  authentication remote pre-share
  authentication local pre-share
  keyring local keyring-ifc-7
  lifetime 3600
  dpd 10 5 periodic
!
crypto ipsec transform-set ikev-transport-ifc-7 esp-gcm 256
  mode tunnel
!
crypto ipsec profile ikev-profile-ifc-7
  set transform-set ikev-transport-ifc-7
  set pfs group14
  set ikev2-profile ikev-profile-ifc-7
  tunnel protection ipsec profile ikev-profile-ifc-7
!
interface Tunnel100
  description To GCP VPN
  vrf forwarding wanVrf
  ip address 169.254.0.14 255.255.255.252
  ip mtu 1400
  ip tcp adjust-mss 1400
  tunnel source GigabitEthernet1
  tunnel mode ipsec ipv4
  tunnel destination 35.220.81.45
  tunnel protection ipsec profile ikev-profile-ifc-7
end
```

The following example shows how to configure BGP:

Example:

```
router bgp 65320
  bgp router-id 172.16.1.1
  bgp log-neighbor-changes
!
address-family ipv4 vrf wanVrf
```



```
network 172.16.8.0 mask 255.255.255.0
network 172.16.9.0 mask 255.255.255.0
redistribute connected
neighbor 169.254.0.9 remote-as 65092
neighbor 169.254.0.9 ebgp-multihop 255
neighbor 169.254.0.9 activate
neighbor 169.254.0.13 remote-as 65092
neighbor 169.254.0.13 ebgp-multihop 255
neighbor 169.254.0.13 activate
exit-address-family
!
```

Step 4 Repeat the previous steps for all external devices.

Configuring Intersite Connectivity Between Google Cloud Site and Other Cloud Sites


The following sections describe how to configure connectivity between two cloud sites. If you want to configure connectivity between a Google Cloud site and an on-premises site, see [Configuring Inter-site Connectivity Between Google Cloud Site and On-Premises Sites, on page 23](#).




Note As Google Cloud only supports non-EVPN connection all cloud sites must be of the same connectivity as Google Cloud which is BGP-IPv4. If other cloud sites use BGP-EVPN, Google Cloud can still be managed, but will not have intersite connectivity to other cloud sites.

Before you begin

- Step 1** Before starting, make sure that the Cloud APIC has hub network configured in atleast one region (at most four region supported) to establish Inter-Site connectivity.
- Step 2** Navigate to the **Site Connectivity**.



Step 3 Choose the site where you want to create the Inter-Site connectivity. Once you choose the site, on the right-hand side window you will see Inter-Site Connectivity and External Connectivity.



Step 4 Under the Inter-Site Connectivity, Choose **Add Site**.

gcpSite44 Settings

0 | 0 | 0 | 0

Inter-Site Connectivity External Connectivity

General ^

APIC Site ID
175

ACI Multi-Site

BGP ^

BGP Autonomous System Number
65112

BGP Password

OSPF ^

Inter-Site Connectivity ⓘ

Site	Protocol	
AzureCapic Connection Type: Public	BGP-IPv4	<input type="checkbox"/>
awsCapic Connection Type: Public	BGP-IPv4	<input type="checkbox"/>

+ Add Site

Connectivity ^

SDA Connectivity ⓘ

Enabled Disabled

Step 5 From the dialog window, under **Connected to Site**, select your cloud APIC site.

Under the Protocol you will see only see BGP-IPv4 as the Connectivity type. This is because you chose Google Cloud site and Google Cloud site only supports BGP-IPv4 connectivity.

AzureCapic → gcpSite44

Please check if CSRs are configured with Public IPs for Public Underlay connection

Connected to Site
gcpSite44

Connection Type
Public Internet

Protocol
Bgpipv4

External VRF *
extVrfAzure

Region	Routers
centralus	ct_routerp_centralus_1
	ct_routerp_centralus_2
	ct_routerp_centralus_0
	ct_routerp_centralus_3

IKE Version
 Version 1
 Version 2

gcpSite44 → AzureCapic

Save

Step 6 Choose the external VRF.

Step 7 From the **External VRF** dropdown, select the external VRF.

This is the external VRF you have configured in [Creating External VRFs in Infra Tenant, on page 22](#).

Version 1
Version 2

gcpSite44 → AzureCapic

i Please check if CSRs are configured with Public IPs for Public Underlay connection

Connected to Site
gcpSite44

Connection Type
Public

Protocol
Bgpipv4

External VRF *
external-vrf1

Region	Cloud Native Router	VPN Router
us-west1	gcphub007	default
us-central1	gcphub007	default

IKE Version
Version 1
Version 2

Save

Step 8 Click **Save** to save the intersite connectivity configuration.

Deploying Infra Configuration

This section describes how to deploy the Infra configuration for the external connectivity from cloud sites.

Step 1 In the top right of the main pane, choose **Deploy > Deploy & Download External Device Config files**.

The **Deploy & Download External Device Config files** option pushes the configuration to the Cloud APIC sites and enables the end-to-end interconnect from the cloud sites to the external devices.

In addition, this option downloads a zip file that contains configuration information that you will use to enable connectivity from external devices to Cloud site router deployed in your cloud sites. A followup screen appears that allows you to select all or some of the configuration files to download.

Step 2 In the confirmation window, click **Yes**.

The Deployment started, refer to left menu for individual site deployment status message will indicate that Infra configuration deployment began and you can verify each site's progress by the icon displayed next to the site's name in the left pane. After successful deployment, you can check the tunnels and BGP sessions created across cloud sites from Cloud APIC dashboard.

Creating an External EPG

This section describes how to create an external EPG in the Infra template using subnet selection. We will use this external EPG to represent the external networks, then configure and apply contracts between the external EPG and the cloud EPG to allow communication between the endpoints in your cloud site and the external networks.

Step 1 In the **Main menu**, select **Application Management > Schemas**.

Step 2 Select the schema and the template that contains your external VRFs.

You can create similar configurations for all (AWS, Azure or Google Cloud) Infra templates, but we recommend using different application profile names in the next step to avoid any possible confusion.

Step 3 Create an **Application Profile** in the template.

You will need to associate the external EPG you create with an application profile.

Step 4 Create and configure an **External EPG**.

- a) Select **Create Object > External EPGs**.
- b) In the external EPG's properties sidebar, select `CLOUD` for **Site Type**.
- c) From the **Application Profile** dropdown, select the profile you created in the previous step.
- d) From the **Virtual Routing and Forwarding** dropdown, select the external VRF you created.

Step 5 Configure the external EPG's site-local properties.

- a) In the left sidebar, select the template under a site to which it is assigned.
- b) In the template's site-local properties, select `External-Site` for **Route Reachability**.
- c) Click **Add Selector**.
- d) In the **Add New Endpoint Selector** dialog, provide the external subnet.

This is an external subnet that requires connectivity to the cloud site, for which you configured route leaking in the previous section. For example, `172.16.8.0/24`.

Step 6 Deploy the templates to create the external EPG in the cloud site.


Importing Google Cloud User Tenant

If you are importing an existing tenant follow the procedure below. If you wish to create a new tenant, refer to this section [Creating Google Cloud User Tenant, on page 38](#).

Step 1 From the Nexus Dashboard's **Service Catalog**, open the Nexus Dashboard Orchestrator service.

You will be automatically logged in using the Nexus Dashboard user's credentials.

Step 2 In the Nexus Dashboard Orchestrator GUI, manage the sites.



- From the left navigation menu, select **Infrastructure** > **Sites**.
- In the main pane, change the **State** from `Unmanaged` to `Managed` for each fabric that you want the Nexus Dashboard Orchestrator to manage.

Step 3 Import the existing cloud tenant.

- In the **Sites** page, click the actions (...) menu next to the site you enabled for management and select **Import Tenants**.
- In the **Import Tenants** dialog, select the tenant you want to import and click **OK**.

Step 4 Verify that the tenant's external connectivity infra configuration was imported successfully.

For external connectivity to be imported, it has to be configured on all the regions in which hub is instantiated.

- Navigate to **Infrastructure** > **Site Connectivity** page.
- Click **Configure**.
- In the **General Settings** page, select the **External Devices** tab.

Verify that the external device is present

- In the **General Settings** page, select the **IPSec Tunnel Subnet Pools** tab.

Verify that the external connectivity subnet pool is present.

- In the left sidebar, select the site from which you imported the tenant.

In the site's settings, select the **External Connectivity** tab and confirm that the external network is present.

Note Do not deploy infra configuration from Nexus Dashboard at this time and proceed to the next section to import the external VRF.

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.

- For either a managed or an unmanaged tenant, you must first set up a project in Google Cloud. See [Setting Up the Google Cloud Project for a User Tenant, on page 36](#) for those instructions.
- For an unmanaged tenant, you must then generate the necessary private key information and download the JSON file from Google Cloud. See [Generating and Downloading Private Key Information from Google Cloud for an Unmanaged Tenant, on page 38](#).

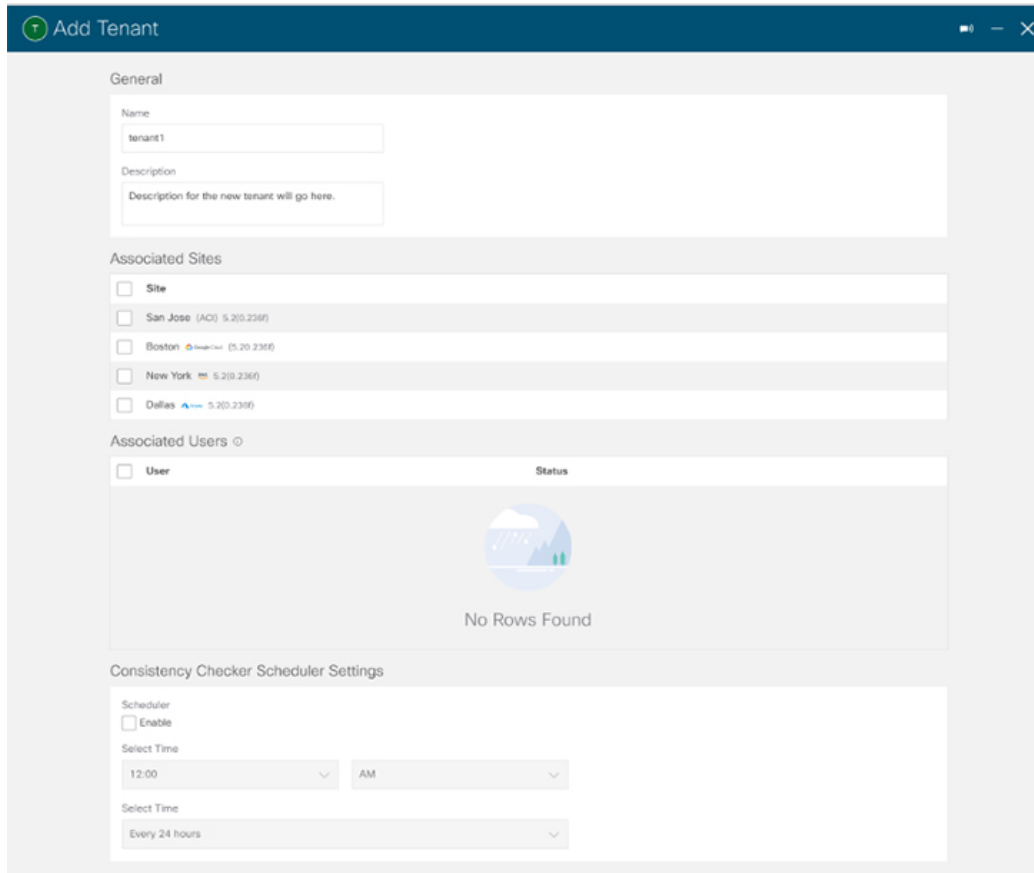
- Step 1** Log in to your Nexus Dashboard Orchestrator.
- Step 2** In the left navigation menu, choose "Tenants".
- Step 3** Choose "Add Tenant".
- Step 4** Under **General**, provide a tenant name and an optional description.

The tenant name must be in the following format:

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

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

- Step 5** From the **Associated Sites** area, choose the Google Cloud site where you want to create the tenant.



The screenshot shows the 'Add Tenant' form with the following details:

- General:** Name: tenant1; Description: Description for the new tenant will go here.
- Associated Sites:**
 - Site
 - San Jose (AQ) 5.20.2360
 - Boston 5.20.2360
 - New York 5.20.2360
 - Dallas 5.20.2360
- Associated Users:** No Rows Found
- Consistency Checker Scheduler Settings:**
 - Scheduler: Enable
 - Select Time: 12:00 AM
 - Select Time: Every 24 hours

- Step 6** After selecting your Google Cloud site, click on the edit icon to specify your account information.

General

Name
tenant1

Description
Description for the new tenant will go here.

Associated Sites

Site	Version
<input type="checkbox"/> Site	
<input type="checkbox"/> San Jose (ACI) 5.2(0.236f)	
<input checked="" type="checkbox"/> Boston (Google Cloud) 5.20.236f	
<input type="checkbox"/> New York 5.2(0.236f)	
<input type="checkbox"/> Dallas 5.2(0.236f)	

Step 7 Fill in all the mandatory information.

General

Security Domains
Select Security Domain(s)

Google Cloud Platform

Google Cloud Project ID *
123456789

Access Type *
Unmanaged Identity Managed Identity

Save

- **Google Cloud Platform ID:** Provide the ID of the Google Cloud user account you have created for this tenant.
- **Access type:** You will have two options under Access type:
 - Choose **Managed Identity** if you want to allow the Cloud APIC VM to manage the cloud resources.
For either a managed or an unmanaged tenant, you must first set up a project in Google Cloud. See [Setting Up the Google Cloud Project for a User Tenant, on page 36](#) for those instructions.
 - Choose **Unmanaged Identity** if you want to manage the cloud resources via a specific application. In this case you must also provide the application's credentials to the Cloud APIC.
 - For either a managed or an unmanaged tenant, you must first set up a project in Google Cloud. See [Setting Up the Google Cloud Project for a User Tenant, on page 36](#) for those instructions.
 - For an unmanaged tenant, you must then generate the necessary private key information and download the JSON file from Google Cloud. See [Generating and Downloading Private Key Information from Google Cloud for an Unmanaged Tenant, on page 38](#).

The **Key Id** and **Client Id** fields appear if you choose **Unmanaged Identity** as the access type.

- **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 38](#).
- **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 38](#).
- **Email:** Enter the email address associated with your Google Cloud project.

Tenant Setting for Boston Cloud Site

General

Security Domains

Name

[Add Security Domain](#)

Google Cloud Platform

Google Cloud Platform ID*

123456789

Access Type*

Unmanaged Identity Managed Identity

Please enter Google Cloud Platform's Service Account Information.

Key ID* Will be visible if Access Type == "Unmanaged"

70b57r48sg890

RSA Private Key

MIIEvAIBADANBgkqhkiG9w0BAQEFAASCbYwggSIAgEAAoIBAQC0Xg3oA011zU15O1ypXCvhy90L...

Client ID*

XYZ

Email*

abc@mail.com

Security Domains for Google Cloud Platform

Name

[Add Security Domain for Google Cloud Platform](#)

Cancel Save

Step 8 Choose **Save** after filling in the configuration for the Google Cloud.

What to do next

If you are creating a managed tenant, you must now set the necessary permissions in Google Cloud for the managed tenant. Go to [Setting the Necessary Permissions in Google Cloud for a Managed Tenant, on page 41](#) 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.
- Click the pencil icon on the row for this service account. The **Edit Permissions** window is displayed.
 - 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.
 - 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
 - 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 Cloud EPGs

We recommend creating cloud objects in a separate template and schema from the Infra tenant configuration (such as external VRFs) you have already done.

Use the following procedure to create a new schema for the Cloud APIC site. For this use-case example, we will configure a single schema and one template.

You are in the Nexus Dashboard Orchestrator for this entire procedure.

-
- Step 1** In the Main menu, click **Schemas**.
- Step 2** On the Schema screen, click the **Add Schema** button.
- Step 3** On the Untitled Schema screen, replace the text `Untitled Schema` at the top of the page with a name for the schema that you intend to create (for example, `schema-1`).
- Step 4** Create a template.
- If your Google cloud site has BGP-EVPN intersite connectivity, choose **ACI Multi-Cloud** template type; if the site has BGP-IPv4 connectivity, choose **Cloud Local**.
- In the left pane, mouse over **Template 1** and click the notepad icon. Then change the template's name, for example in Google Cloud case `template1-gcp`.
 - In the middle pane, click the area **To build your schema please click here to select a tenant**.
 - In the right pane, access the **Select A Tenant** dialog box and choose the tenant you want. This is the tenant you imported [Importing Google Cloud User Tenant, on page 34](#) or created in [Creating Google Cloud User Tenant, on page 38](#).
- Step 5** After choosing the tenant, create an **Application Profile** in the template.
- You will need to associate the cloud EPG you create with an application profile.
- Step 6** Create and configure a **Cloud EPG**.
- Select **Create Object > Cloud EPGs**.
 - From the **Application Profile** dropdown, select the profile you created in the previous step.
 - From the **Virtual Routing and Forwarding** dropdown, select the cloud VRF you created.
 - In the right-hand properties sidebar, select the cloud VRF you created for this EPG.
- Step 7** Assign the template you just created to the Google Cloud site.
- Step 8** Configure the cloud EPG's site-local properties.
- In the left sidebar, select the template under a site to which it is assigned.
 - In the template's site-local properties, select `Cloud Site` for **Route Reachability**.

Creating Schema, Template and VRFs for your Google Cloud Site

- Step 1** In the Main menu, click **Schemas**.
- Step 2** On the Schema screen, click the **Add Schema** button.
- Step 3** On the Untitled Schema screen, replace the text `Untitled Schema` at the top of the page with a name for the schema that you intend to create (for example, `schema-1`).
- Step 4** Configure the first template.
- If your Google cloud site has BGP-EVPN intersite connectivity, choose **ACI Multi-Cloud** template type; if the site has BGP-IPv4 connectivity, choose **Cloud Local**.

- Step 5** In the left pane, mouse over **Template 1** and click the notepad icon. Then change the template's name (for example, `template1-gcp`).
- Step 6** Navigate to your cloud template.
- Step 7** Choose **Add VRF** under VRFs, then enter the display name and description for the VRF.
- Step 8** Click on the VRF that you just created.
The Template Properties and Site Local Properties are displayed on the right side of your screen.
- Step 9** Under Site Level Properties, choose **Add Region**.
In the pop-up, select the region that you want.
- Step 10** After selecting the region, choose **Add CIDR**.
Enter the CIDR information for the VRF.
- Choose **Primary** if you are adding a primary CIDR.
 - Choose **Secondary** if you are adding a secondary CIDR.
- Step 11** Enter the Subnet and Subnet Group Label.
When creating a subnet, you will use the **Subnet Group Label** to assign a unique label to a specific subnet group. For more details on configuring CIDR, subnets, and subnet group labels, see "Understanding VPCs and Subnets Under Google Cloud and Cloud Context Profiles Under Cloud APIC" in the [Cisco Cloud APIC for Google Cloud User Guide](#).
- Step 12** Choose **Save**.

Configuring an Application Profile and EPG

- Step 1** In the middle pane click + **Application Profile**.
- Step 2** In the right pane, enter the Application Profile name in the **Display Name** field (for example, `app1`).
- Step 3** In the middle pane, click + **Add EPG**.
- Step 4** In the right pane, enter an EPG name in the **Display Name** field.
- Step 5** In the **Cloud Properties** area, you can see the VRF you just created in the previous section (for example, `cloud-vrf`).
-

Adding Cloud Endpoint Selector

On the Cloud APIC, 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 CIDR. You define the endpoints for a cloud EPG using an object called endpoint selector. The endpoint selector is essentially a set of rules run against the cloud instances assigned to either AWS, Azure or Google Cloud managed by the Cloud APIC. Any endpoint selector rules that match endpoint instances will assign that endpoint to the Cloud EPG. Unlike the traditional on-premises ACI fabrics where endpoints can only belong to a single EPG at any one time, it is possible to configure endpoint selectors to match multiple Cloud EPGs. This in turn would cause the same instance to belong to multiple Cloud EPGs. However, we recommend configuring endpoint selectors in such a way that

each endpoint matches only a single EPG. The section below will walk you through the process of Adding End point Selector.

- Step 1** In the Nexus Dashboard Orchestrator, select the EPG you create in the previous section.
- Step 2** In the right pane, in the **Site Local Properties** area, click + **Selector** under the Selectors heading to configure the endpoint selector.

If you plan to stretch this EPG, you can also choose to add the endpoint selector at the template level instead.

- Step 3** In the **Add New End Point Selector** form, enter a name in the **End Point Selector Name** field, based on the classification that you use for this endpoint selector.

For example, for an endpoint selector with the IP Subnet classification, you might use a name such as `IP-Subnet-EPSelector`.

- Step 4** Click + **Expression**, then use the three fields to configure the endpoint selector based on how you want to classify the endpoints in the cloud:

The **Type** field determines the expression that you want to use for the endpoint selector:

- Choose **IP Address** if you want to use an individual IP address or a subnet for the endpoint selector.
- Choose **Region** if you want to use the cloud region for the endpoint selector, then choose the specific region that you want use.

When you select `Region` for the endpoint selector, every instance within the tenant that is brought up in that region will be assigned to this cloud EPG.
- Choose **Custom tags or labels** if you want to create a custom tag or label for the endpoint selector. Start typing to enter the custom tag or label, then click **Create** on the new field to create a new custom tab or label.

The **Operator** field determines the relation between the type and its value:

- **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 if the expression contains only a key.

The **Value** field determines the collection of endpoints that you want to use for the endpoint selector, based on the choices that you made for the two previous fields. This can be a single IP address, a subnet, AWS region or zone, or a custom tag value.

For this use case, you will be assigning endpoints based on IP subnets, so you will configure the endpoint selector using the following example values:

- **Type**: `IP Address`
- **Operator**: `Equals`
- **Value**: `3.3.1.0/24`

Step 5 Click the checkmark next to the new endpoint selector.

Step 6 Click **Save** in the Add New End Point Selector form.

Applying Contract Between External EPG and Cloud EPG

This section describes how to apply a contract to allow communication between the endpoints in your cloud site and the external networks. One thing to keep in mind regarding Google Cloud contracts is that the contracts should be deployed bi-directionally for bi-directional traffic.

Before you begin

- You must have one or more cloud EPGs [Creating Cloud EPGs, on page 42](#) already configured in your cloud site.
- You must have external EPG [Creating an External EPG, on page 34](#) and external VRF [Creating External VRFs in Infra Tenant, on page 22](#) already configured.

Step 1 In the **Main menu**, select **Application Management > Schemas**.

Step 2 Create a contract and assign it to the cloud EPG.

- Select the schema and the template that contains your existing cloud EPG.
- Create the contract you will use for this use case.

If you already have an existing contract you want to apply for communication between the external network and the Cloud EPG, you can skip this step.

Otherwise, create a contract and the required filters as you typically would for any inter-EPG communication in Cisco ACI fabrics.

- Assign the contract to the cloud EPG.

You can decide which of the two EPGs (cloud EPG and external EPG) will be the `provider` and which will be the `consumer` based on your specific use case.

Step 3 Assign the contract to the External EPG.

- Select the schema and the template where you created your external EPG.
- Assign the contract to the external EPG.

If you configured your cloud EPG to be the provider, choose `consumer` for the external EPG; otherwise, if the cloud EPG is the consumer, choose `provider`.

Note Contract scope must be set to "global", so the contract can be used between external EPG and cloud EPG.

Step 4 Deploy the templates.

Configuring Route Leaking Between Cloud VRF and External VRF

This use case focuses on route leaking between a google cloud VRF (in a user tenant) and an external VRF (in Infra tenant) in order to establish traffic flow between your google cloud site and another cloud or on-premises site. If you want to configure route leaking between two cloud VRFs (for example, to enable traffic flow within the same Google Cloud site), see [Configuring Route Leaking between Two Cloud VRFs, on page 60](#).

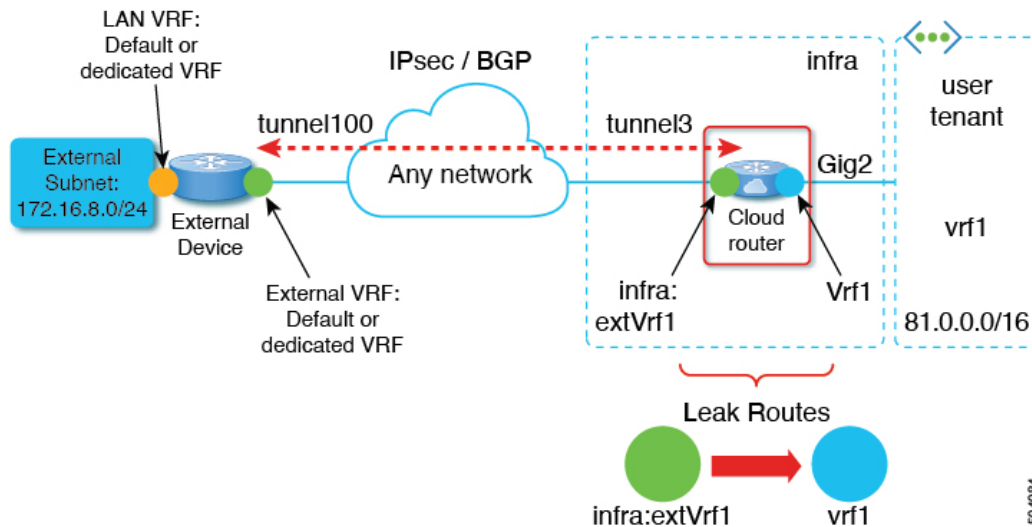
Before you begin

You must have one or more cloud VRFs already configured in your cloud site. You will configure route leaking from the external VRF to an existing cloud VRF.

Step 1 In the **Main menu**, select **Application Management > Schemas**.

Step 2 Configure route leaking from external VRF to a cloud VRF.

The following steps show how to configure the following route leaking:



- Open the schema where you created the Infra tenant template containing the external VRF.
- In the left sidebar under **SITES**, select that specific template associated to the cloud site.
- In the site-local properties, select the external VRF defined in the template.

This is the VRF you created and assigned to one or more external devices .

- In the VRF's right-hand properties sidebar, click **+Add Leak Route**.

The **Add Leak Routes** dialog will open.

- In the **Add Leak Routes** dialog's settings area, click **Select a VRF** and choose a cloud VRF.

The goal of this step is to leak routes from the external VRF to the cloud VRFs, so select the cloud VRF to which you want to leak routes from the external VRF whose properties you are configuring.



CHAPTER 5

Configuring Internal Connectivity for Google Cloud Workloads

- [Internal Connectivity Workflow](#), on page 49
- [Importing Google Cloud User Tenant](#), on page 49
- [Creating a Tenant](#), on page 50
- [Creating Schema, Template and VRFs for your Google Cloud Site](#), on page 58
- [Creating Cloud EPGs](#), on page 58
- [Applying contract between the cloud EPGs](#), on page 59
- [Configuring Route Leaking between Two Cloud VRFs](#), on page 60

Internal Connectivity Workflow

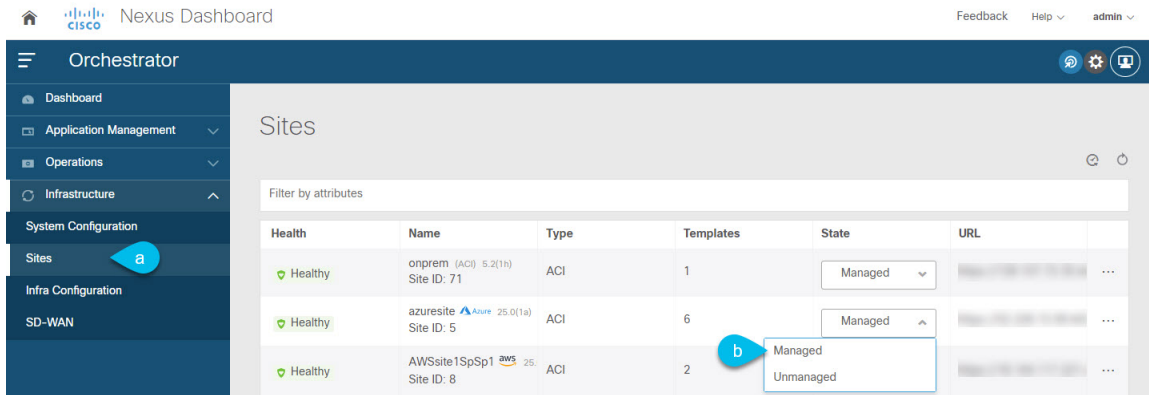
The following sections describe how to configure Google Cloud sites infra, intersite connectivity, and a simple deployment use case. The workflow includes:

- Select the EPG you create in the previous section
- Configuring route leaking between cloud VRFs
- Creating or importing a Google cloud user tenant and EPGs and applying contracts to enable communication between sites

Importing Google Cloud User Tenant

If you are importing an existing tenant follow the procedure below. If you wish to create a new tenant, refer to this section [Creating Google Cloud User Tenant](#), on page 38.

-
- Step 1** From the Nexus Dashboard's **Service Catalog**, open the Nexus Dashboard Orchestrator service. You will be automatically logged in using the Nexus Dashboard user's credentials.
- Step 2** In the Nexus Dashboard Orchestrator GUI, manage the sites.



- From the left navigation menu, select **Infrastructure** > **Sites**.
- In the main pane, change the **State** from `Unmanaged` to `Managed` for each fabric that you want the Nexus Dashboard Orchestrator to manage.

Step 3 Import the existing cloud tenant.

- In the **Sites** page, click the actions (...) menu next to the site you enabled for management and select **Import Tenants**.
- In the **Import Tenants** dialog, select the tenant you want to import and click **OK**.

Step 4 Verify that the tenant's external connectivity infra configuration was imported successfully.

For external connectivity to be imported, it has to be configured on all the regions in which hub is instantiated.

- Navigate to **Infrastructure** > **Site Connectivity** page.
- Click **Configure**.
- In the **General Settings** page, select the **External Devices** tab.
Verify that the external device is present
- In the **General Settings** page, select the **IPSec Tunnel Subnet Pools** tab.
Verify that the external connectivity subnet pool is present.
- In the left sidebar, select the site from which you imported the tenant.
In the site's settings, select the **External Connectivity** tab and confirm that the external network is present.

Note Do not deploy infra configuration from Nexus Dashboard at this time and proceed to the next section to import the external VRF.

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.

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.

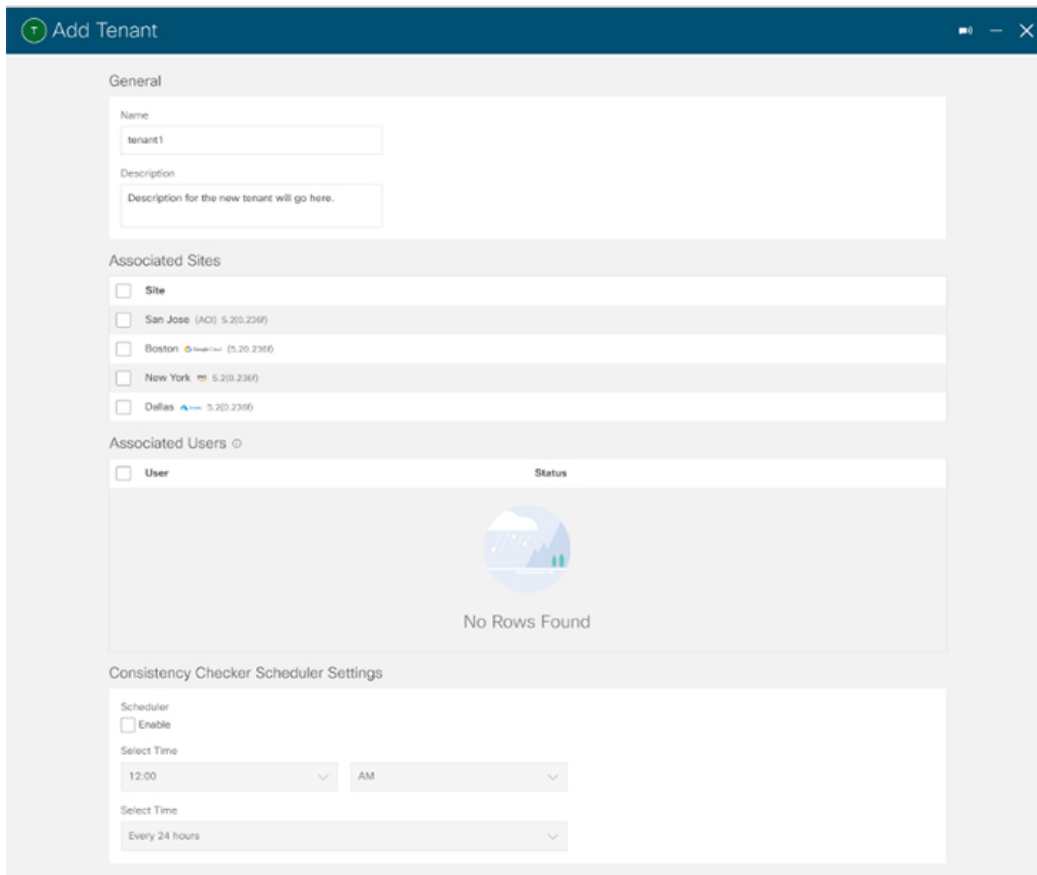
- Step 1** Log in to your Nexus Dashboard Orchestrator.
- Step 2** In the left navigation menu, choose "Tenants".
- Step 3** Choose "Add Tenant".
- Step 4** Under **General**, provide a tenant name and an optional description.

The tenant name must be in the following format:

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

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

- Step 5** From the **Associated Sites** area, choose the Google Cloud site where you want to create the tenant.



The screenshot shows the 'Add Tenant' form with the following details:

- General:** Name: tenant1; Description: Description for the new tenant will go here.
- Associated Sites:**

Site
<input type="checkbox"/> San Jose (AC) 5.2(0.236f)
<input type="checkbox"/> Boston 5.2(0.236f)
<input type="checkbox"/> New York 5.2(0.236f)
<input type="checkbox"/> Dallas 5.2(0.236f)
- Associated Users:** No Rows Found
- Consistency Checker Scheduler Settings:**
 - Scheduler: Enable
 - Select Time: 12:00 AM
 - Select Time: Every 24 hours

- Step 6** After selecting your Google Cloud site, click on the edit icon to specify your account information.

Step 7 Fill in all the mandatory information.

- **Google Cloud Platform ID:** Provide the ID of the Google Cloud user account you have created for this tenant.
- **Access type:** You will have two options under Access type:
 - Choose **Managed Identity** if you want to allow the Cloud APIC VM to manage the cloud resources.
For either a managed or an unmanaged tenant, you must first set up a project in Google Cloud. See [Setting Up the Google Cloud Project for a User Tenant, on page 36](#) for those instructions.
 - Choose **Unmanaged Identity** if you want to manage the cloud resources via a specific application. In this case you must also provide the application's credentials to the Cloud APIC.
 - For either a managed or an unmanaged tenant, you must first set up a project in Google Cloud. See [Setting Up the Google Cloud Project for a User Tenant, on page 36](#) for those instructions.
 - For an unmanaged tenant, you must then generate the necessary private key information and download the JSON file from Google Cloud. See [Generating and Downloading Private Key Information from Google Cloud for an Unmanaged Tenant, on page 38](#).

The **Key Id** and **Client Id** fields appear if you choose **Unmanaged Identity** as the access type.

- **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 38.
- **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 38.
- **Email:** Enter the email address associated with your Google Cloud project.

Tenant Setting for Boston Cloud Site

General

Security Domains

Name

[Add Security Domain](#)

Google Cloud Platform

Google Cloud Platform ID*

123456789

Access Type*

Unmanaged Identity Managed Identity

Please enter Google Cloud Platform's Service Account Information.

Key ID* Will be visible if Access Type == "Unmanaged"

70b67148og890

RSA Private Key

MIIEvABADANlJgkqkG0w0BAQEFAASCBywggSIAgEAoIBAQCOXg3oAQ11ZU1501ygXCvhy9CL...

Client ID*

XYZ

Email*

abc@mail.com

Security Domains for Google Cloud Platform

Name

[Add Security Domain for Google Cloud Platform](#)

Cancel Save

Step 8 Choose **Save** after filling in the configuration for the Google Cloud.

What to do next

If you are creating a managed tenant, you must now set the necessary permissions in Google Cloud for the managed tenant. Go to [Setting the Necessary Permissions in Google Cloud for a Managed Tenant](#), on page 41 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.
- Click the pencil icon on the row for this service account. The **Edit Permissions** window is displayed.
 - 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.
 - 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
 - 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 Schema, Template and VRFs for your Google Cloud Site

- Step 1** In the Main menu, click **Schemas**.
- Step 2** On the Schema screen, click the **Add Schema** button.
- Step 3** On the Untitled Schema screen, replace the text `Untitled Schema` at the top of the page with a name for the schema that you intend to create (for example, `schema-1`).
- Step 4** Configure the first template.
If your Google cloud site has BGP-EVPN intersite connectivity, choose **ACI Multi-Cloud** template type; if the site has BGP-IPv4 connectivity, choose **Cloud Local**.
- Step 5** In the left pane, mouse over **Template 1** and click the notepad icon. Then change the template's name (for example, `template1-gcp`).
- Step 6** Navigate to your cloud template.
- Step 7** Choose **Add VRF** under VRFs, then enter the display name and description for the VRF.
- Step 8** Click on the VRF that you just created.
The Template Properties and Site Local Properties are displayed on the right side of your screen.
- Step 9** Under Site Level Properties, choose **Add Region**.
In the pop-up, select the region that you want.
- Step 10** After selecting the region, choose **Add CIDR**.
Enter the CIDR information for the VRF.
- Choose **Primary** if you are adding a primary CIDR.
 - Choose **Secondary** if you are adding a secondary CIDR.
- Step 11** Enter the Subnet and Subnet Group Label.
When creating a subnet, you will use the **Subnet Group Label** to assign a unique label to a specific subnet group. For more details on configuring CIDR, subnets, and subnet group labels, see "Understanding VPCs and Subnets Under Google Cloud and Cloud Context Profiles Under Cloud APIC" in the [Cisco Cloud APIC for Google Cloud User Guide](#).
- Step 12** Choose **Save**.
-

Creating Cloud EPGs

We recommend creating cloud objects in a separate template and schema from the Infra tenant configuration (such as external VRFs) you have already done.

Use the following procedure to create a new schema for the Cloud APIC site. For this use-case example, we will configure a single schema and one template.

You are in the Nexus Dashboard Orchestrator for this entire procedure.

-
- Step 1** In the Main menu, click **Schemas**.
- Step 2** On the Schema screen, click the **Add Schema** button.
- Step 3** On the Untitled Schema screen, replace the text `Untitled Schema` at the top of the page with a name for the schema that you intend to create (for example, `schema-1`).
- Step 4** Create a template.
- If your Google cloud site has BGP-EVPN intersite connectivity, choose **ACI Multi-Cloud** template type; if the site has BGP-IPv4 connectivity, choose **Cloud Local**.
- In the left pane, mouse over **Template 1** and click the notepad icon. Then change the template's name, for example in Google Cloud case `template1-gcp`.
 - In the middle pane, click the area **To build your schema please click here to select a tenant**.
 - In the right pane, access the **Select A Tenant** dialog box and choose the tenant you want. This is the tenant you imported [Importing Google Cloud User Tenant, on page 34](#) or created in [Creating Google Cloud User Tenant, on page 38](#).
- Step 5** After choosing the tenant, create an **Application Profile** in the template.
- You will need to associate the cloud EPG you create with an application profile.
- Step 6** Create and configure a **Cloud EPG**.
- Select **Create Object > Cloud EPGs**.
 - From the **Application Profile** dropdown, select the profile you created in the previous step.
 - From the **Virtual Routing and Forwarding** dropdown, select the cloud VRF you created.
 - In the right-hand properties sidebar, select the cloud VRF you created for this EPG.
- Step 7** Assign the template you just created to the Google Cloud site.
- Step 8** Configure the cloud EPG's site-local properties.
- In the left sidebar, select the template under a site to which it is assigned.
 - In the template's site-local properties, select `Cloud Site` for **Route Reachability**.

Applying contract between the cloud EPGs

This section describes how to apply a contract to allow communication between the endpoints with in your cloud site. One thing to keep in mind regarding Google Cloud contracts is that the contracts should be deployed bi-directionally for bi-directional traffic.

Before you begin

You must have multiple cloud EPGs [Creating Cloud EPGs, on page 42](#) already configured in your cloud site.

-
- Step 1** In the **Main menu**, select **Application Management > Schemas**.
- Step 2** Create a contract and assign it to the cloud EPG.
- Select the schema and the template that contains your existing cloud EPG.
 - Create the contract you will use for this use case.

If you already have an existing contract you want to apply for communication between the Cloud EPGs, you can skip this step.

Otherwise, create a contract and the required filters as you typically would for any inter-EPG communication in Cisco ACI fabrics.

- c) Assign the contract to the cloud EPG.

You can decide which of the two EPGs will be the `provider` and which will be the `consumer` based on your specific use case.

Step 3 Select the other EPG.

- a) From the right property side bar , choose **Add contract**.
- b) In the contract window, select which contract you want to assign.
- c) Select the same contract you assigned in previous step.
- d) Click **Save**

Step 4 Deploy the templates.

Configuring Route Leaking between Two Cloud VRFs

This use case focuses on route leaking between two internal cloud VRFs. You must have multiple cloud VRFs already configured in your cloud site. If you want to configure route leaking between a cloud VRF an external VRF (for example, to enable external connectivity for your Google Cloud site to another site), see [Configuring Route Leaking Between Cloud VRF and External VRF, on page 47](#)

Step 1 In the **Main menu**, select **Application Management > Schemas**.

Step 2 Configure route leaking from Cloud VRF-1 to a cloud VRF-2.

The following steps show how to configure the following route leaking:

- a) Open the schema where you created the Infra tenant template containing the first cloud VRF.
- b) In the left sidebar under **SITES**, select that specific template associated to the cloud site.
- c) In the site-local properties, select the cloud VRF defined in the template.
- d) In the VRF's right-hand properties sidebar, click **+Add Leak Route**.

The **Add Leak Routes** dialog will open.

- e) In the **Add Leak Routes** dialog's settings area, click **Select a VRF** and choose a cloud VRF.
- f) In the **Add Leak Routes** dialog, choose **Leak All** routes.

After selecting **Leak All**, the subnet IP will be populated with `0.0.0.0/0` to leak all routes.

- g) Click **Save** to save the route leak configuration.
- h) Select the template and click **Deploy** to deploy the configuration.

Step 3 Configure route leaking from a cloud VRF-2 to the cloud VRF-1.

- a) Open the schema which contains the template that defines your cloud VRF.
- b) In the left sidebar under **SITES**, select the specific cloud site.

- c) In the site-local properties, select the cloud VRF.
- d) In the VRF's right-hand properties sidebar, click **+Add Leak Route**.

The **Add Leak Routes** dialog will open.

- e) In the **Add Leak Routes** dialog's settings area, click **Select a VRF** and choose the internal VRF.

The goal of this step is to leak routes between the cloud VRFs

- f) In the **Add Leak Routes** dialog, choose **Leak All** routes.
 - g) Click **Save** to save the route leak configuration.
 - h) Select the template and click **Deploy** to deploy the configuration.
-

