Cisco SD-WAN Cloud onRamp for Multicloud using Google Cloud Platform

Prescriptive Deployment Guide

Nov, 2021

© 2020 Cisco and/or its affiliates. All rights reserved.

Contents

© 2020 Cisco and/or its affiliates. All rights reserved.

Introduction

About the Guide

This document discusses the design and deployment of Cisco SD-WAN Cloud onRamp for Multicloud using Google Cloud Platform (GCP). This guide focuses on the design and configuration of the Site to Google Cloud connectivity and the site-to-site connectivity through Google global network. The document includes some of the best practices and steps to instantiate a pair of Cisco Catalyst 8000v instances within a Google cloud gateway(s), association of Google Host VPCs within tags, establishment of Intra-tag communication, mapping of tags to service side VPN along with necessary design and steps to allow the site-to-site communication via the Google Global Network.

This guide assumes that the Cisco SD-WAN controllers are already deployed and integrated into the vManage NMS controller, the WAN Edge devices are deployed, and the Cisco SD-WAN overlay network is successfully established. This guide is written based on the configuration supported in 20.5/17.5 release. Refer to Appendix B to view the device models and software versions used in this deployment and refer to Appendix C for the Catalyst 8000v CLI configuration used in the deployment

This document contains four major sections:

- The Define section introduces the Cisco Cloud onRamp for Multi-cloud feature and explains the overall solution, along with the benefits of deploying it.
- The Design section includes the two use cases covered in the guide, along with the design components and considerations for the successful integration of Cisco SD-WAN and Google Cloud.
- The Deploy section includes all the prerequisites and the necessary steps to associate Google cloud with vManage NMS, along with the steps to deploy vManage device templates for the Catalyst8000v devices to be hosted in Google cloud. The section also includes the steps for the automated deployment of pair of Cisco Catalyst 8000V instances in cloud gateways with their interfaces anchored in three different VPCs to support the two use cases presented within the design section.
- The Operate section explains some of the common monitoring and troubleshooting capabilities available within the Cisco vManage for the Cisco Cloud onRamp for Multi-Cloud feature.

Figure 1. **Implementation Flow**

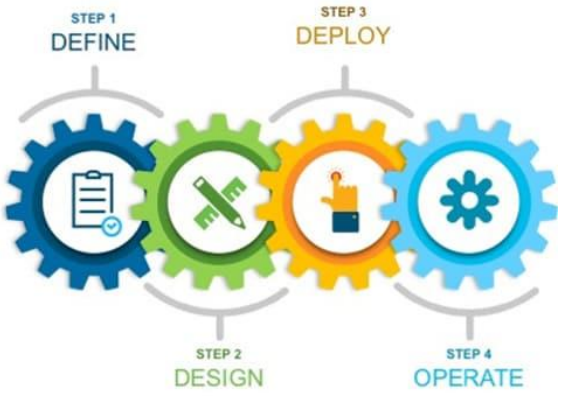

© 2020 Cisco and/or its affiliates. All rights reserved.

Audience

The audience for this document includes network design engineers, network operations personnel, and cloud operations personnel who wish to establish access from branch site to a service hosted in a VPC in Google Cloud or to connect branches across different regions through the GCP global network.

Define - Cisco Cloud onRamp for Multi-Cloud Introduction

Cisco SD-WAN Interconnection with Multi-Cloud

Network engineers in today's world of enterprise IT are beginning to understand the benefits of multicloud fabric for it offers a premium experience when connecting your branch network to workloads and SaaS applications. Enterprise customers now start to adopt more than one cloud. They choose different cloud platforms to deploy different services based on what works best for them. For example, the front-end service could be placed in AWS, while SQL service is provided by Azure and analytics service from Google Cloud Platform (GCP).

The rest of this guide explains some of the Cisco SD-WAN Google cloud connectivity models using Cisco SD-WAN cloud hub. For details regarding Cisco SD-WAN AWS interconnection, refer to the Cisco Cloud OnRamp for IaaS using AWS guide, and for details regarding Cisco SD-WAN Azure interconnection, refer to the Cisco Cloud OnRamp for IaaS using Azure guide.

About the Solution - Cisco SD-WAN Interconnection with Google Cloud

The Cisco SD-WAN Google cloud integration enables network policies, such as Cisco SD-WAN segmentation, to follow network traffic across the boundary between the enterprise network and Google Cloud, for end-to-end control of security, performance, and quality of experience. With this new integration, customers can extend a single point of orchestration and management for their Cisco SD-WAN network to include the underlay offered by Google Cloud backbone.

Using the Cisco Cloud OnRamp for Multi-Cloud feature, you can associate your Google cloud account with Cisco vManage to bring up a pair of redundant Cisco Catalyst 8000V Edge Software (Cisco Catalyst 8000V) in a Google Cloud Gateway. The services hosted in the Google VPCs are discovered and mapped to the service-side VPNs of the Catalyst 8000v Cisco SD-WAN edge devices. Therefore, allowing your on-premises/ cloud hosted SD-WAN sites to access services hosted within the Google platform. This type of Google cloud connectivity model is referred to as Site-to-Cloud (S2C) connectivity. To enable this form of mapping, a site-to-cloud VPC is automatically provisioned in Google cloud using the Cisco multi-cloud solution.

The Cisco Cloud OnRamp for Multi-Cloud solution also enables you to establish communication between two on-premises branches using your Google platform as the underlay. To enable this traffic flow across two sites through Google cloud, you need to have at least two cloud gateways brought up in different Google cloud regions, each containing a pair of Catalyst8000v devices. A site-to-site VPC (S2S VPC) is automatically brought up in each of the Catalyst 8000v devices hosted between the two separate cloud gateways. The Site-to-Site VPC contains the subnets and firewalls associated with the second WAN tunnel interface labeled as private1. IPsec or GRE tunnels are brought up between the Site-to-Site VPC (S2S VPC) in each Catalyst8000v devices located in the cloud gateways, therefore allowing branch networks located in different regions to communicate with each other through this automated WAN tunnel in GCP Global Network. This type of Google cloud connectivity model is referred to as Site-to-Site (S2S) connectivity.

© 2020 Cisco and/or its affiliates. All rights reserved.

Within the Google Cloud Platform (GCP), the following services are used to interconnect the google resources with the on-premises Cisco SD-WAN network.

Table 1. Google Terminology

Google Terminology	Definition
Google Virtual Private Cloud (VPC) 	Google VPC is similar to a traditional network that you operate in your own data center, except that its virtualized within Google Cloud. It is a global resource that consists of a list of regional virtual subnetworks so that each VPC network is logically isolated from each other within the google cloud. Each VPC network implements the following: A virtual firewall to control incoming and outgoing traffic directly on the hypervisor, routes to govern the traffic leaving VM instances/ VPC network and forwarding rules to direct traffic to a google cloud resource in a VPC network.
Service Account Keys 	A service account key is essentially a public/private key pair. IAM roles applied against the service account key define the privileges a service account key has. The private key associated to your service account is required to associate your Google Account with vManage. The steps to create private key is explained in the deploy section of this guide.
Network Connectivity Center (NCC) 	Network Connectivity Center is a hub-and-spoke model for network connectivity management in Google Cloud, that enables you to connect your on-premises SD-WAN branch and data center networks together by using Google's network as a network for data transfer.
Google Cloud Router 	Cloud Router is a fully distributed and managed Google Cloud service that programs custom dynamic routes and scales with your network traffic. The Google cloud routers take care of all the underlay routing within the Google domain.
Google Billing ID 	Google Billing ID is a unique number assigned to the organization or an individual that the account is registered to.

Note, all the Google IDs needed for the programmatic sign-in access are explained in depth in the Deploy section, along with the necessary steps.

Benefits of Deploying Cisco Cloud onRamp for Multi-Cloud using Google Cloud Platform

Some of the key benefits of deploying the Cloud onRamp feature includes the following:

Table 2. Benefits of deploying cloud OnRamp for multi-cloud solution

Automated Infrastructure in Public Cloud (Extend SD-WAN)	Cisco SD-WAN Cloud onRamp securely extends Cisco SD-WAN's fabric to public cloud environments through a simplified and automated process. By using Cisco SD-WAN Cloud onRamp, customers can reduce the overall time to deploy and connect branch offices to cloud workloads in minutes. It helps enterprises to significantly increase productivity and avoid error-prone, manual processes.
Policy Control (Policy Framework)	Users can fully utilize Cisco SD-WAN capabilities in the cloud. This allows all on-premises data centers and branch locations which are part of the Cisco SD-WAN fabric to leverage features available both on Cisco SD-WAN vManage and within the cloud provider. Some of the features that can be leveraged on Cisco SD-WAN fabric include:

© 2020 Cisco and/or its affiliates. All rights reserved.

	• Using centralized control policy to redirect critical site to site traffic to traverse through the Google cloud platform, while routing non-critical traffic through SD-WAN overlay network. • Embedded security features such as IPS/IDS, Stateful Firewall, AMP, and URL Filtering to protect and filter data traffic as it leaves the on-premises network to access the Internet cloud.
Reduce OPEX (Cost Effective)	Expenses shift from fixed costs for hardware, software, and data center infrastructure to variable costs based on the usage of compute resources available on Google public cloud.

Design - Cisco Cloud onRamp for Multi-Cloud Use Case and Feature Overview

The design section is organized in the following order:

- Cloud Hub Use Case #1 – Site-to-Cloud Connectivity (S2C)
- Cloud Hub Use Case #2 – Site-to-Site Connectivity (S2S) with SDWAN Core
- Design Components and Consideration
  - Supported Platform and Software
  - Preparation of Device Templates
  - Creation and Management of Cloud Gateway
  - Mapping of VPN to Tag
  - Site-to-Cloud Workflow
  - Site-to-Site Workflow

Cloud Hub Use Case #1 - Site-to-Cloud Connectivity (S2C)

In the site-to-cloud use case, Cisco Cloud onRamp automates the backend processes to host a pair of Catalyst 8000v devices as a part of the Cloud Gateway (CGW), through which on-premises branch and data center Cisco SD-WAN Edge devices access Google cloud resources and applications running in a VPC. In this scenario, a branch site, such as site 1 connects to a VPN 0 transport tunnel interface part of the catalyst 8000v device, which connects to your applications hosted in Google VPCs through the Catalyst 8000v service side VPN interface.

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 2. **Site-to-Cloud Model**

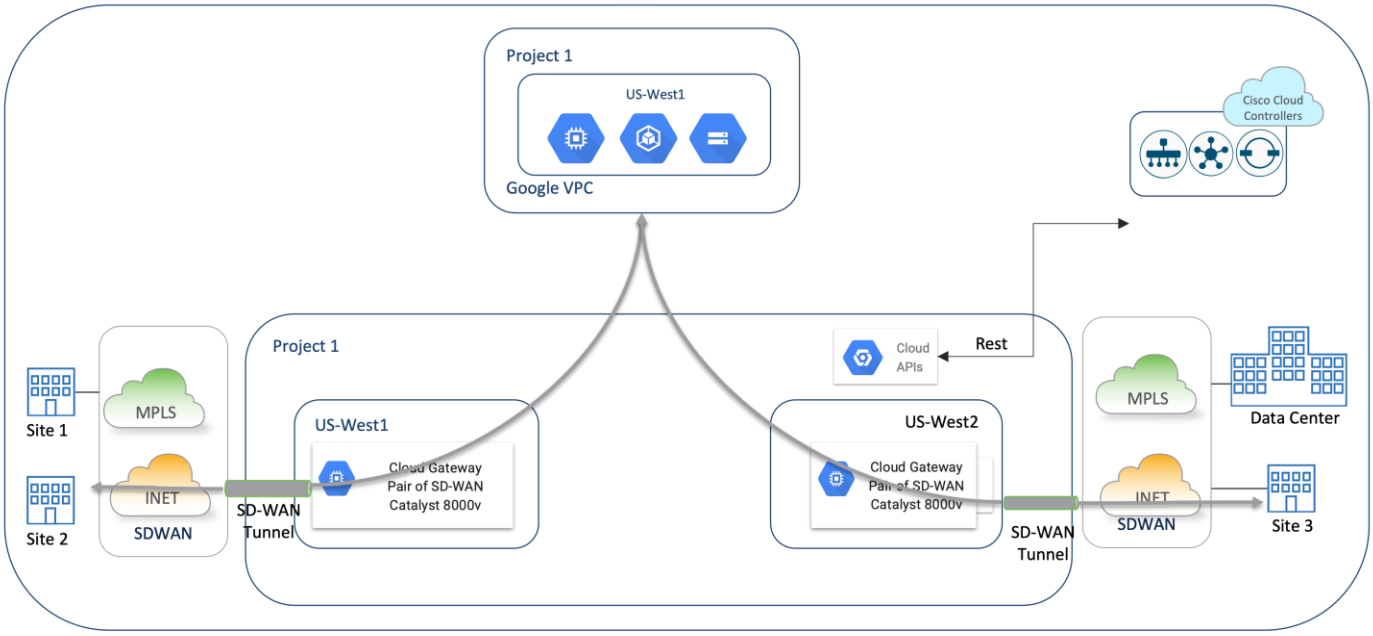

Cloud Hub Use Case #2 – Site to Site Connectivity (S2S) with SDWAN Core

To achieve the site-to-site connectivity, two branches located in different Google Cloud regions are connected through the site-to-site transit VPC or Cloud Gateway in Google Global Network. Within the Google cloud platform, a pair of Catalyst 8000v Edge devices are deployed in a Cloud Gateway (CGW) with a 2nd WAN transport interface in the site-to-site VPC. Similarly, a second set of Catalyst 8000v edge devices are deployed in another Cloud Gateway (CGW) with its 2nd WAN transport interface also part of the same site-to-site VPC. This second transport interface can be either IPsec or GRE encapsulated. The two Cloud Gateways (CGWs) are placed in different Google cloud regions and the Catalyst 8000v devices in each of the CGWs appear as spokes to the Network Connectivity (NCC) hub. The NCC hub acts as a transit within the Google domain to carry data traffic across the IPsec or GRE tunnel via the global Google backbone.

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 3. **Site-to-Site Model**

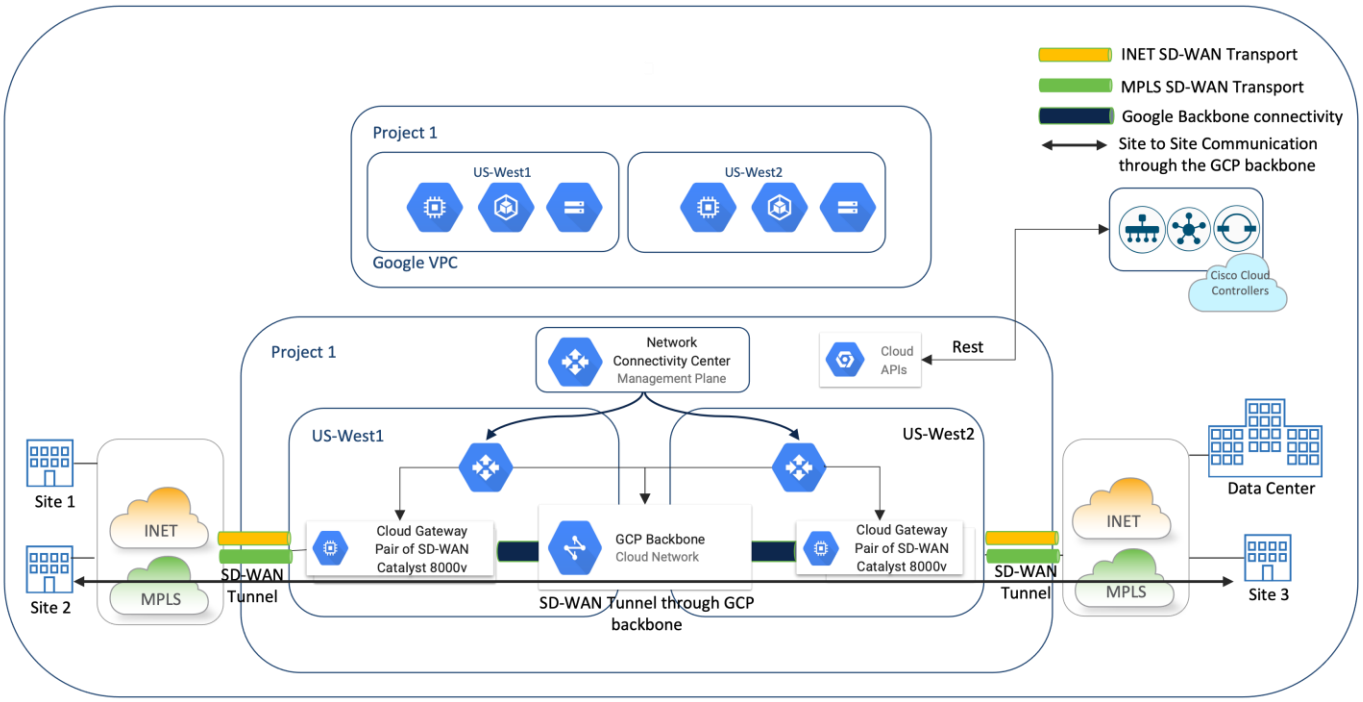

Design Components and Consideration

The main building blocks of the Site-to-Cloud (S2C) and Site-to-Site (S2S) design include the bring up of two Cisco Catalyst 8000v SD-WAN routers in each of the Google Cloud Gateways (CGWs). Each of the Catalyst 8000v device contained within these gateways have three gigabit ethernet interfaces. Each interface in the Catalyst 8000v is a part of a different VPC, therefore the design includes the bring up and configuration of the three VPCs – WAN VPC, Site-to-Cloud (S2C) VPC and Site-to-Site (S2S) VPC. Each of these VPCs are explained in detail within the design section.

The design section also focuses on the supported cloud platform/ software and features to be considered while deploying the SD-WAN Google Cloud Integration using Cisco SD-WAN Cloud onRamp for Multi-Cloud feature.

Supported Platforms and Software

Cloud onRamp for Multi-cloud using Google Cloud is supported on Catalyst 8000v SD-WAN router running IOS-XE SD-WAN version 17.5 and above, with the vManage controller running version 20.5 or a higher code.

Preparation of Device Templates

Having noted the software and hardware requirements, the next steps include the design and management of a device template that comprises of a list of feature templates to be associated with a pair of unused Catalyst8000v devices.

To deploy just the Site-to-Cloud use case, you need only a pair of Catalyst 8000v devices deployed in a single Cloud Gateway (CGW).

To deploy the Site-to-Cloud and the Site-to-Site use case, you need at least a two pairs of Catalyst 8000v devices deployed in two different Cloud Gateways (CGWs). The two CGWs containing a pair of Catalyst 8000vs must be deployed in two different Google cloud regions.

© 2020 Cisco and/or its affiliates. All rights reserved.

For ease of use, the GCP device templates are pre-built in the vManage NMS under the default templates tab. The default device template includes all the necessary feature templates required to bring up the WAN Edge devices with both control plane and data plane connections established. While one transport interface is provisioned using the default device template in WAN VPC, the service side VPN interface in the Site-to-Cloud (S2C) VPC and the 2nd WAN transport interface in the Site-to-Site (S2S) VPC is deployed automatically using the Cloud onRamp workflow.

It is upon the creation and deployment of Cloud Gateways (CGWs) in Google cloud that the interfaces and its associated resources required to complete the site-to-cloud and site-to-site use case are deployed.

Creation and Management of Cloud Gateway

A Cloud Gateway (CGW) is simply a pair of Catalyst 8000v devices provisioned in a Google Cloud region. Within the Google cloud platform, CGWs are instantiated in different Google cloud regions based on the global or custom settings entered within the vManage Cloud onRamp setup page.

During the instantiation of a Cloud Gateway containing a pair of Catalyst8000v devices, as mentioned previously, three Google VPCs are created:

WAN VPC: The WAN interface is a part of the WAN VPC. The WAN VPN interface template for the Catalyst 8000v devices is created based on the device template attached to it. Its interface IP address is assigned based on the subnet entered within the CGW global settings. Similarly, the location of the VPC is based on the location in which the CGW is deployed.

Within the Catalyst 8000v, the Gigabit Ethernet 1 or Gig1 interface is part of the WAN VPC and note this is the only interface in the Catalyst 8000v that is assigned a public IP address.

Site-to-Cloud (S2C) VPC: A second VPC is created for the site to cloud use case and associated to the service side VPN interface. The VPC, VPN and the VPN interface details are automatically created and configured based on the region and subnet details entered within the CGW global and custom settings.

The peering between the S2C VPC and host VPC is done based on the BGP offset details entered in the global settings page.

Note, within the Catalyst 8000v, the Gigabit Ethernet 2 or Gig2 interface is part of the S2C VPC.

Site-to-Site (S2S) VPC: And finally, a third VPC is created on enabling the site-to-site connectivity within the CGWs global settings. This VPC is associated with a second WAN interface labeled color private1. The VPC, VPN and the VPN interface details are automatically created and configured region and subnet details entered within the CGW global and custom settings.

The BGP routing protocol is used to learn and advertise the underlay routes.

Note, within the Catalyst 8000v, the Gigabit Ethernet 3 or Gig3 interface is part of the S2S VPC.

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 4. **VPCs created in GCP**

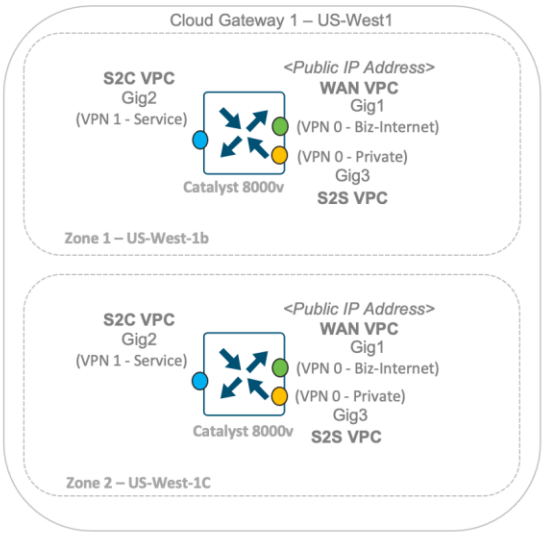

Cloud Gateway Settings

The rest of this section focuses on understanding the Cloud Gateway (CGW) settings required to bring up a pair of Catalyst 8000v devices within the CGWs.

Catalyst 8000v Instance Size: The Catalyst v8000 WAN Edge virtual devices provisioned in google cloud VPC can be of instance type N1-standard-8 and N1-standard-4. Falling under the general-purpose machine category this machine type provides predefined vCPU and memory resources for an instance. The N1-standard-8 machine type provides 8 virtual CPUs and a total of 30GB memory, while the N1-standard-4 machine type provides 4 virtual CPUs and a total of 15GB memory.

Region: Plan the regions wherein your CGWs and its associated resources are to be instantiated. The following link lists out all the regions you can choose to bring up your cloud resources:
https://cloud.google.com/compute/docs/regions-zones

Regions have three or more zones. For example, the us-west1 region denotes a region on the west coast of the United States that has three zones: us-west1-a, us-west1-b, and us-west1-c. During the instantiation of a Cloud Gateway, a pair of Catalyst 8000v devices running 17.5 IOS-XE SD-WAN code are automatically placed in separate zones for redundancy within the chosen region.

CGW Subnet Pool: The subnet pool entered at the time of provisioning your cloud gateways needs to be between the subnet range of /16 to /21. By default, the system allocates /27 per VPC, therefore plan your IP addressing accordingly.

For instance, if the entered subnet pool is 10.52.0.0/16, and you have two Cloud Gateways deployed, one in US-West1 and the other in US-West2 with site-to-site connectivity enabled, then the VPC subnets are assigned subnets in the following order:

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 5. **Subnets in S2S, S2C and WAN VPC**

s2c-enb-solutions-21615		2	1460	Custom	
	us-west1	s2c-enb-solutions-21615-subnet-0			10.52.0.32/27
	us-west2	s2c-enb-solutions-21615-subnet-0			10.52.0.160/27
s2s-enb-solutions-21615		2	1460	Custom	
	us-west1	s2s-enb-solutions-21615-subnet-0			10.52.0.64/27
	us-west2	s2s-enb-solutions-21615-subnet-0			10.52.0.192/27
wan-enb-solutions-21615		2	1460	Custom	
	us-west1	wan-enb-solutions-21615-subnet-0			10.52.0.0/27
	us-west2	wan-enb-solutions-21615-subnet-0			10.52.0.128/27

Based on the subnets listed in the table above, the IP addresses and subnets are assigned to the interfaces part of two Catalyst 8000v devices that belong to the same Cloud Gateway placed in US-West1 in the following way,

Figure 6. **Subnet and IP address distribution in a CGW**

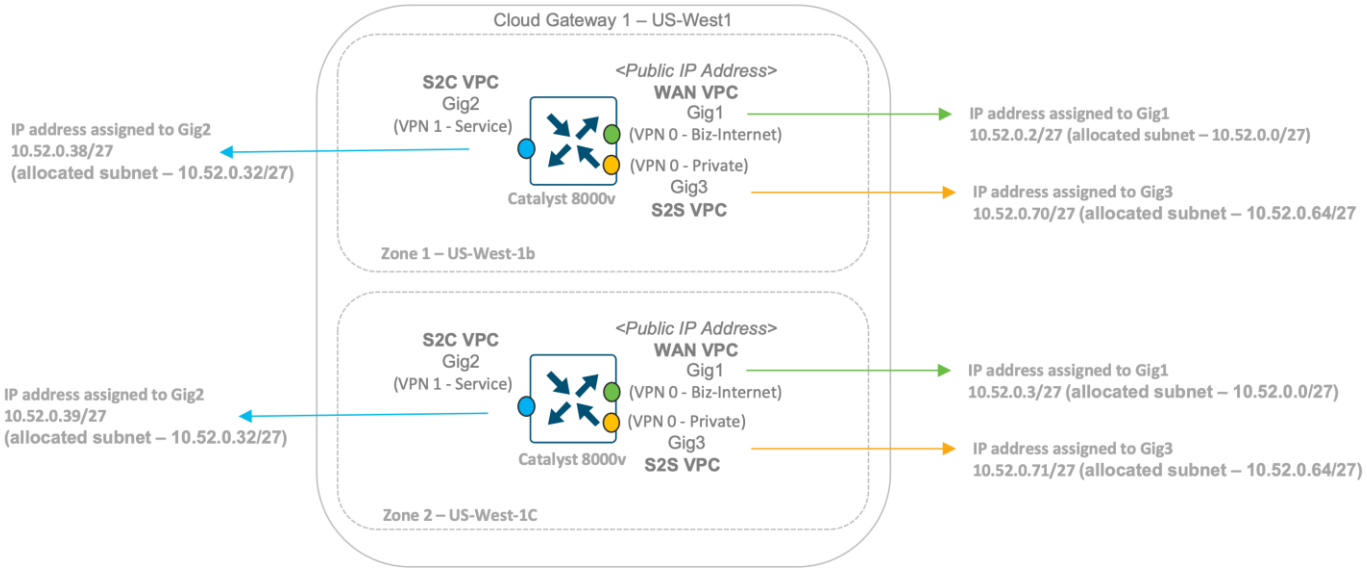

The WAN Interface is part of the WAN VPC, wan-enb-solutions-21615 and this interface is assigned one private and one public IP address. The public IP address is used to establish control connections with the Cisco SD-WAN controllers and data plane connections with branch sites.

To establish the site-to-cloud use case, a single service-side VPN is automatically deployed within each Catalyst 8000v. The service-side VPN (VRF 1/ VPN 1) is deployed and associated to the Site-to-Cloud VPC, s2c-enb-solutions-21615.

To establish the site-to-site use case, a second WAN tunnel interface is brought up under the s2s-enb-solutions-21615 VPC.

BGP ASN Offset: As mentioned earlier, the BGP routing protocol is used to learn and advertise the underlay routes.

To establish this peering BGP ASN offset is entered within the settings page. The BGP ASNs entered in the settings are assigned to the resources as given below:

The entered BGP ASN Offset is the number assigned as the BGP ASN for the GCR in Site-to-Cloud VPC.

© 2020 Cisco and/or its affiliates. All rights reserved.

The entered BGP ASN Offset plus one (BGP ASN + 1) is the ASN assigned for GCRs in Site-to-Site VPC.

The entered BGP ASN Offset plus 10 (BGP ASN +10) will be the ASN assigned to the Catalyst8000v devices.

All ASNs must belong to private ASN space i.e. between the ASN range 64250 to 65520. Plan your BGP ASN numbers accordingly.

Intra-Tag Communication: Within the Cloud onRamp workflow, the discovered VPCs are added to a Tag. Tagging allows for the grouping of several VPCs together to treat them as a single unit.

Within the cloud architecture, communication within the same Google VPC network conversation is possible, but to enable communication between multiple VPC's, you can combine them all under one tag and this establishes intra-tag VPN communication. Therefore, this allows Google VPCs that are a part of the same tag to communicate with other. So, in order to achieve intra tag communication between Google VPC's, links are created between different VPC's. The number of links will be N x (N-1) where N is the number of VPC's.

For example, if you have 4 VPCs, VPC A, VPC B, VPC C and VPC D hosted in Google cloud, and they were all added to the same tag A, then the number of links between the VPCs will be, 4 X (4-1) = 12.

Figure 7. **Tagging Operation**

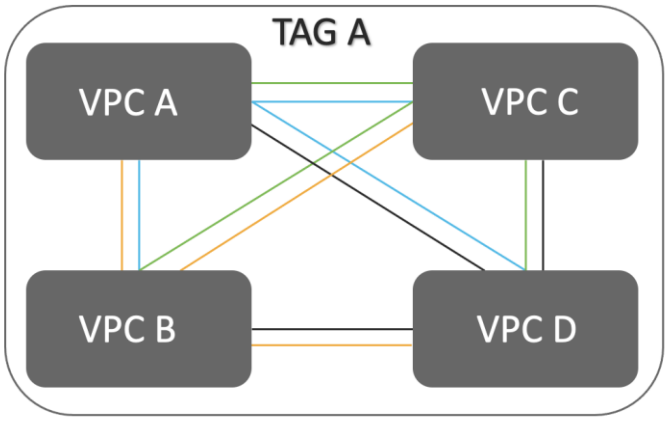

The intra tag communication is enabled by default, however you can choose to disable this from the global settings.

In this guide, customer hosted VPCs named Host VPC-1 and Host VPC-3 are part of the same tag, Tag 1 and customer hosted VPCs named Host VPC-2 and Host VPC-4 are part of another tag, Tag 2.

Figure 8. **Host VPC to Tag mapping**

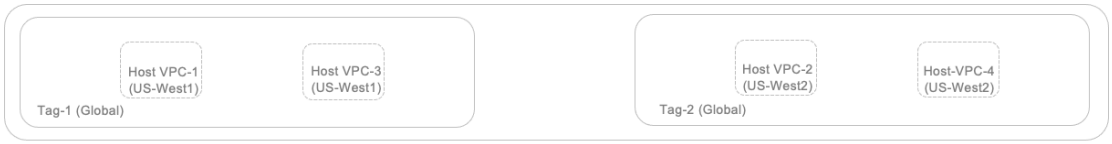

Technical Tip
A VPC hosted in Google cloud platform can be a part of tag or removed from an associated tag or moved from one tag to another tag as part of editing of tags. The design for the tagging operation is the same across all cloud types, AWS and Azure included.

© 2020 Cisco and/or its affiliates. All rights reserved.

Site-to-Site Connectivity and Tunnel Encapsulation Type: Enable site-to-site connectivity to deploy the site-to-site use case.

For the site-to-site use case, you need at least two cloud gateways deployed in two different Google cloud regions. Note, only a single Google Cloud Gateway is supported per Google region. Therefore, to establish site-to-site connectivity you need to bring up two Google Cloud Gateway, placed in two different Google cloud regions.

On enabling site-to-site connectivity within each of the cloud gateways, an additional WAN tunnel interface is automatically created and added through the Cloud onRamp workflow with its color automatically set to private1. This interface is associated with the Site-to -Site VPC (s2s-enb-solutions-21615) and is assigned a private IP address. The tunnel type can be of type GRE or IPsec encapsulated and, regardless of the encapsulation type chosen, all data traffic traversing through the Google backbone is encrypted.

Note: All the site-to-site traffic flows through the Google global network only via the private1 color WAN tunnel interface and not through any other WAN interface created on the Catalyst 8000v devices.

Technical Tip
Google recommends using GRE encapsulated tunnels as it provides twice the bandwidth/throughput.

Network Service Tier: Choose one of the GCP network service tiers – Premium or Standard. Premium: Gives users exceptional high performing network experience by using Google's global network costs, while still delivering performance comparable with other cloud providers. The default is Premium. When updated, it will be effective for the next CGW instantiation.

Mapping of VPN to Tag

Within the cloud onRamp intent matrix page, you can map a service-side VPN to a tag, or map a tag to tag, or map a tag to a VPN. This type of mapping is referred to as a Forward Mapping.

Since the intent is global in nature, if the mapping intent was entered earlier for a CGW, then when a new CGW gets instantiated in a different region the applicable VPN to Tag mapping gets utilized. So, one does not have to retrigger the mapping on the instantiation of a new CGW. Any mapping operation that comes into play at that time is known as Derived Mapping.

Technical Tip

© 2020 Cisco and/or its affiliates. All rights reserved.

In the Cloud onRamp workflow, only one tag can be mapped to one Cisco SD-WAN Service side VPN. For example, if VPN 1 to a tag mapping already exists, then another VPN 1 to tag mapping cannot be completed.

Also, note the intent matrix page lists all the service side VPNs known to vManage from the VPN segments page. So, even though your Catalyst 8000v is configured with only one service-side VPN, the intent matrix mapping page may list 1 or more VPNs.

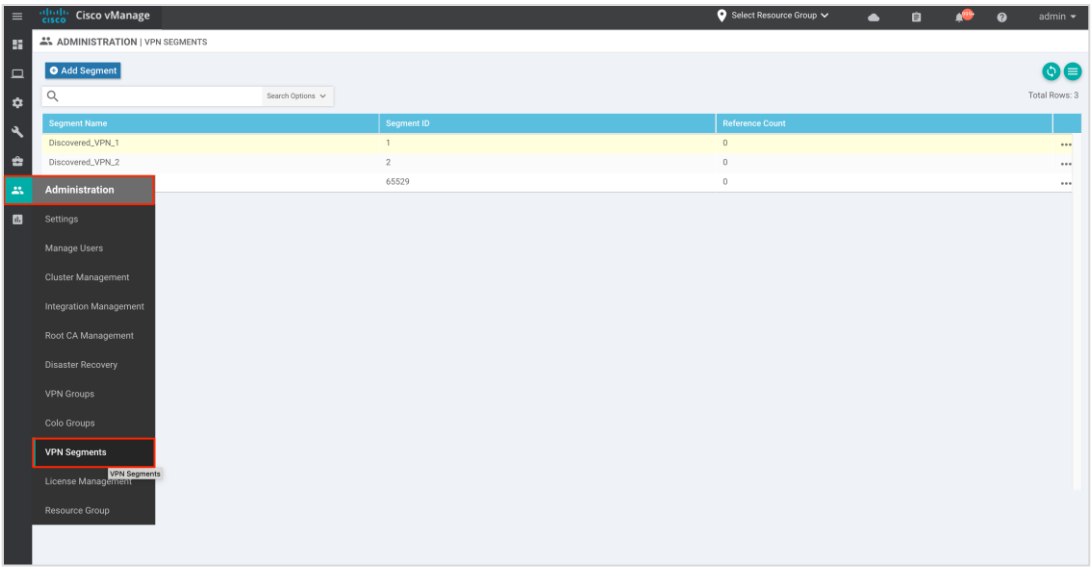

Site-to-Cloud (S2C) VPN to Tag Mapping and Traffic Flow

In the following example figure, two CGWs are deployed, one in US-West 1 and the other in the US-West 2 region with each CGW containing a pair of Catalyst 8000v configured with one transport interface in WAN VPC, and one service side interface in Site to cloud VPC.

Now, during the creation of a CGW, an NCC Site-to-Cloud (S2C) hub is automatically created and each Catalyst 8000v router in the CGW connects as a spoke to the Site-to-Cloud (S2C) Hub. The S2C Hub has a global scope while the Cisco SD-WAN CGW Catalyst8000v devices are instantiated at regional level within the project. While the NCC hubs are essentially responsible for inter-region routing using Google's backbone, the S2C Hub does not participate in any inter-region transit and practically not used in any part of the packet transit.

Therefore, the following figure shows two Cloud gateways, each containing a pair of Catalyst 8000v devices that act as spokes to the Site to Cloud NCC Hub. The VPCs placed on the topmost part of the picture are the customer hosted VPCs in Google cloud. In this guide, customer hosted VPCs named Host VPC-1 and Host VPC-3 are part of the same tag, Tag 1 and customer hosted VPCs named Host VPC-2 and Host VPC-4 are part of another tag, Tag 2.

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 9. **Site to Cloud Design**

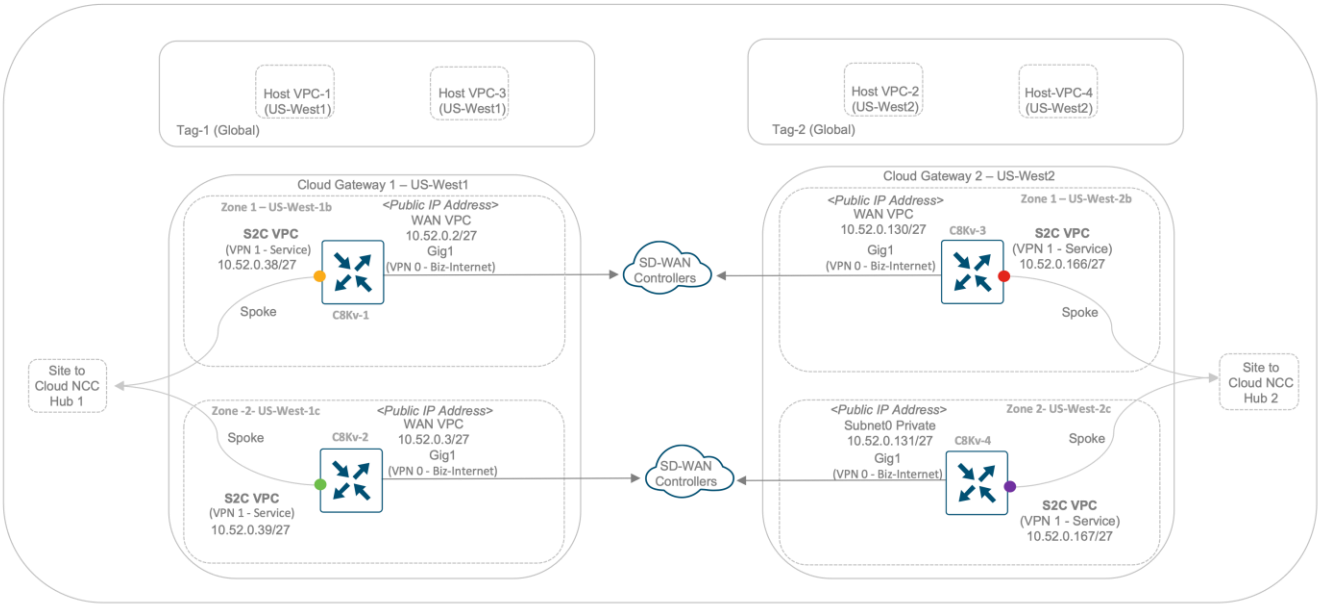

The next steps include the mapping of the service side VPN, VPN 1 to the tag, therefore establishing the communication between branch sites with the VPCs part of the global Tag. As a part of this VPN-tag mapping, Google cloud routers are configured in the Google cloud platform within the Site to Cloud VPC which carries the underlay traffic across.

Since we have two Catalyst 8000v devices deployed in each Cloud gateway (total of four Catalyst 8000v routers), an equivalent number of redundant GCRs are also deployed.

In US-West1, two GCRs are deployed with each GCR assigned two private IP addresses. A pair of GCRs located in a Google cloud region are assigned a total of 4 private IP addresses and the BGP ASN offset entered within the global settings is assigned to each of the GCRs in the site to cloud VPC.

Note: The GCR's are assigned IP addresses based on the IP subnet entered globally or per CGW.

In this guide, the entered BGP ASN Offset is 64520, therefore this BGP ASN Offset is assigned to the GCRs and note each of these GCRs share a common gateway IP address. Note, this gateway IP address is automatically assigned to the GCRs based on the entered subnet. The IP addresses for the first set of the GCRs in the Site to Cloud VPC, range from 10.52.0.34 – 10.52.0.37, therefore IP address 10.52.0.33 is automatically assigned as the gateway IP. Do note, all Google cloud routers are just a conceptual router in the google cloud. The purpose of Cloud Router is to dynamically exchange the routes between two Virtual Private Cloud Networks or Virtual Private Cloud Network and on-premises networks. Cloud Router uses "Border Gateway Protocol" to exchange routing information between the networks. GCRs are not virtual machines or devices that you can log into via telnet or SSH. Similarly, is the gateway IP address too. It is conceptual in nature, and therefore the gateway IP address is not captured in the figures displayed within this guide.

Similarly, part of the tagging operation two redundant GCRs are also automatically provisioned in US-West2 region, with each GCR in that region also assigned a pair of private IP addresses and a BGP ASN (Assigned ASN 64520). This GCR pair also shares a common gateway IP address. The IP addresses for the second set of the GCRs in the Site to Cloud VPC range from 10.52.0.162 – 10.52.0.165, therefore IP address 10.52.0.161 is automatically assigned as the gateway IP.

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 10. **Google Cloud Routers in S2C VPC**

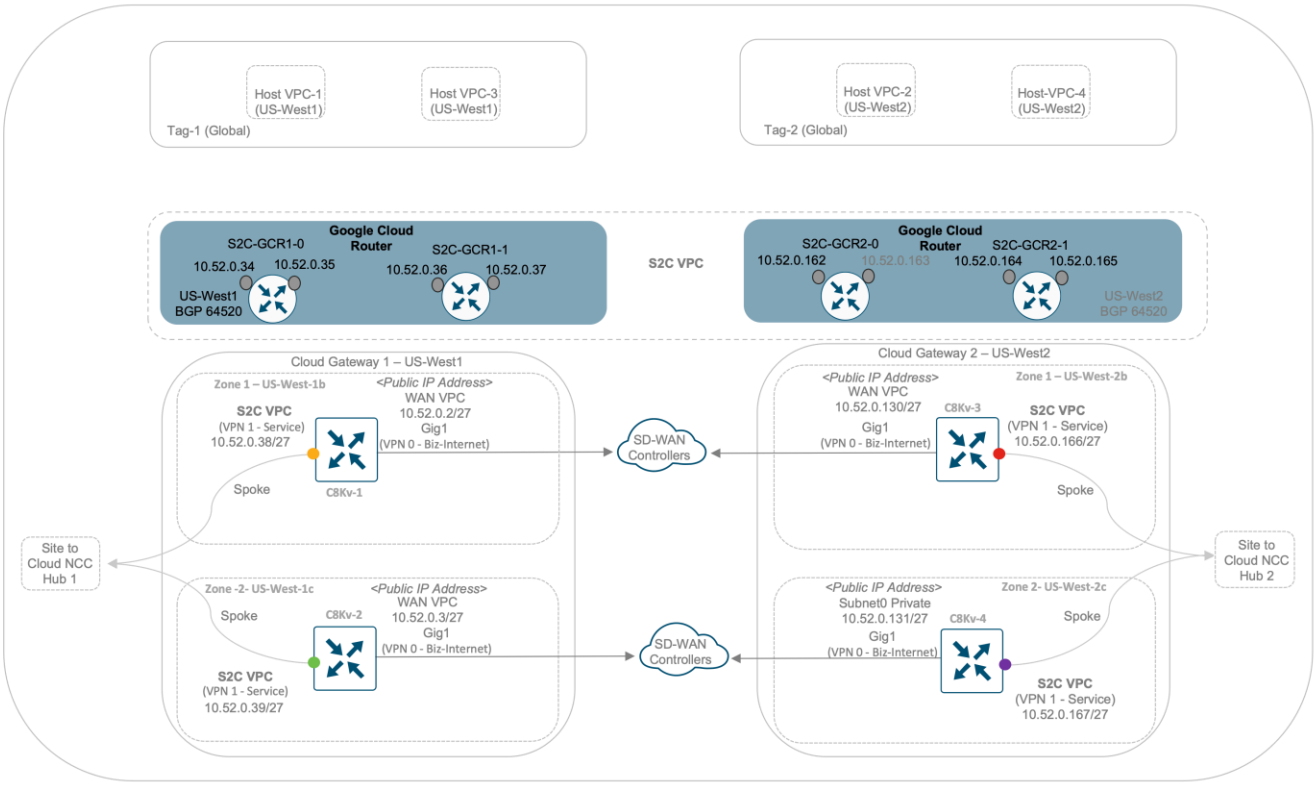

On mapping a service side VPN 1 or VRF 1 S2C interface to a tag, the Catalyst 8000v devices in each CGW peer with the Google Cloud Routers via its Gig2 interface. As explained in the example figure below, on mapping VPN 1 interface to tags, both the Cat8kv-1 and Cat8kv-2 located in US-West1 forms BGP peers with the Google Cloud Routers (GCR) virtual routers via its service-side private IP addresses and similarly, both the Catalyst8000v routers (Cat8kv-3 and Cat8kv-4) in Cloud Gateway 2 located in US-West2 also peers with the GCRs via its service side VPN 1 private IP addresses - 10.52.0.39 and 10.52.0.167.

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 11. **Site-to-Cloud Workflow**

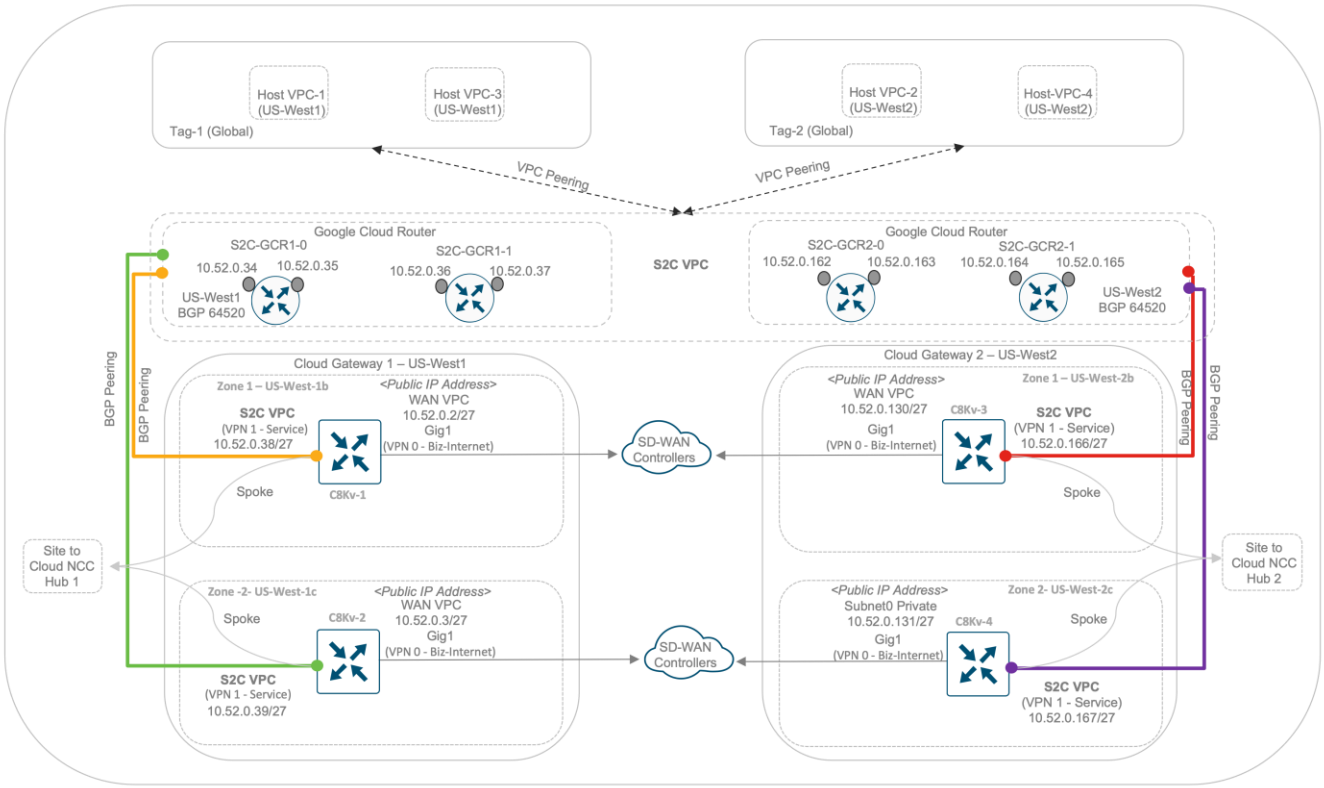

All Google Cloud Routers (GCRs) act as the underlay routing engine in the Google domain and therefore all Catalyst 8000v routers in the CGWs form four eBGP peers with the GCR IP addresses. The example figure below focuses on the eBGP peering between one of the Catalyst8000v routers and a pair of GCRs. The GCRs advertise the host VPC subnets to the Cat8kv-1 device via eBGP.

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 12. **BGP Peering in S2C Workflow**

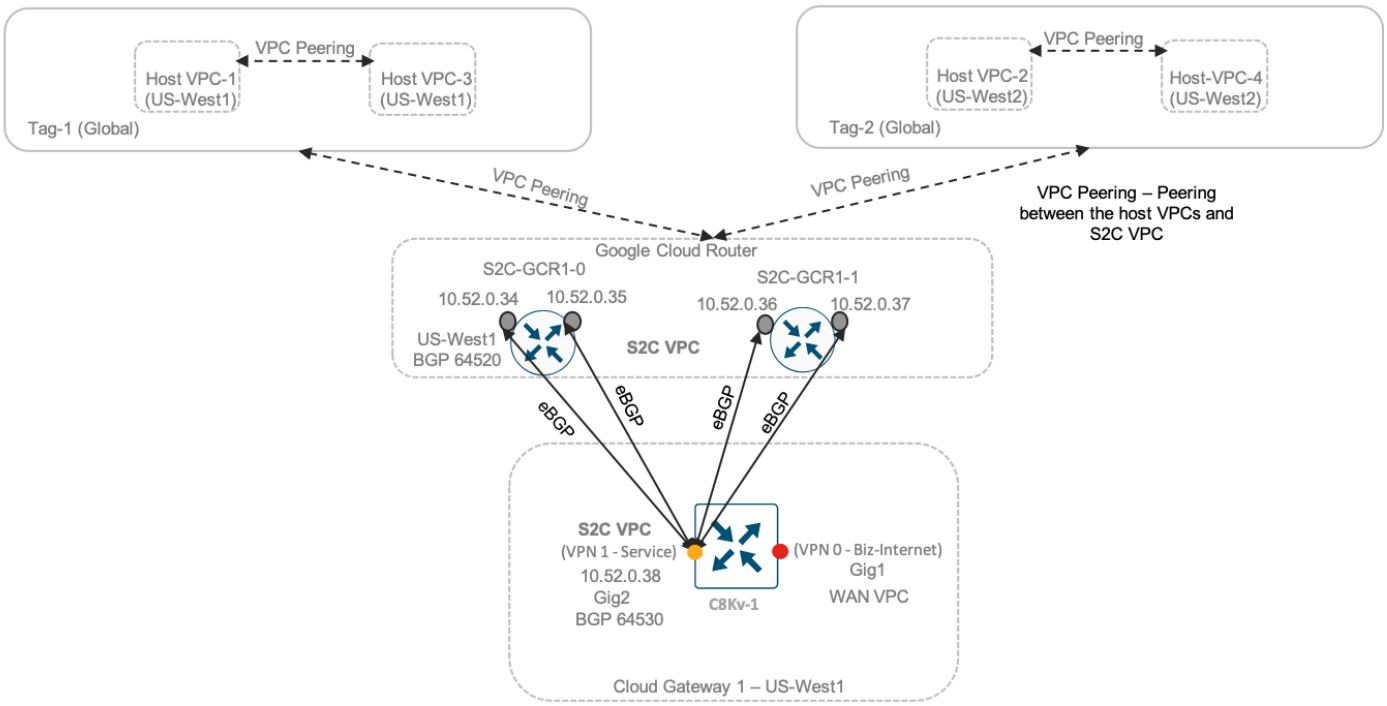

Similarly, each of the Catalyst 8000v devices in each of the Cloud Gateway (CGWs) form 4 eBGP peers with the GCRs (peer with IP 10.52.0.34, 10.52.0.35, 10.52.0.36, 10.52.0.37). Within each Catalyst 8000v device, the learnt BGP routes are advertised to the vSmart controller via OMP. The vSmart controller sends the advertised OMP routes to other Cisco SD-WAN Edge devices. The Catalyst 8000v devices installs the other GCR subnets learnt via OMP and redistributes this into BGP.

The following figure and outputs include the BGP routes learnt by Cat8kv-1. The Cat8kv-1 device installs the second GCR subnet 10.52.0.160/27 learnt via OMP and redistributes this into BGP. For all the relevant CLI outputs from all Catalyst8000v devices, refer to the operate section of this guide.

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 13. **BGP Peering for two Catalyst8000v Devices**

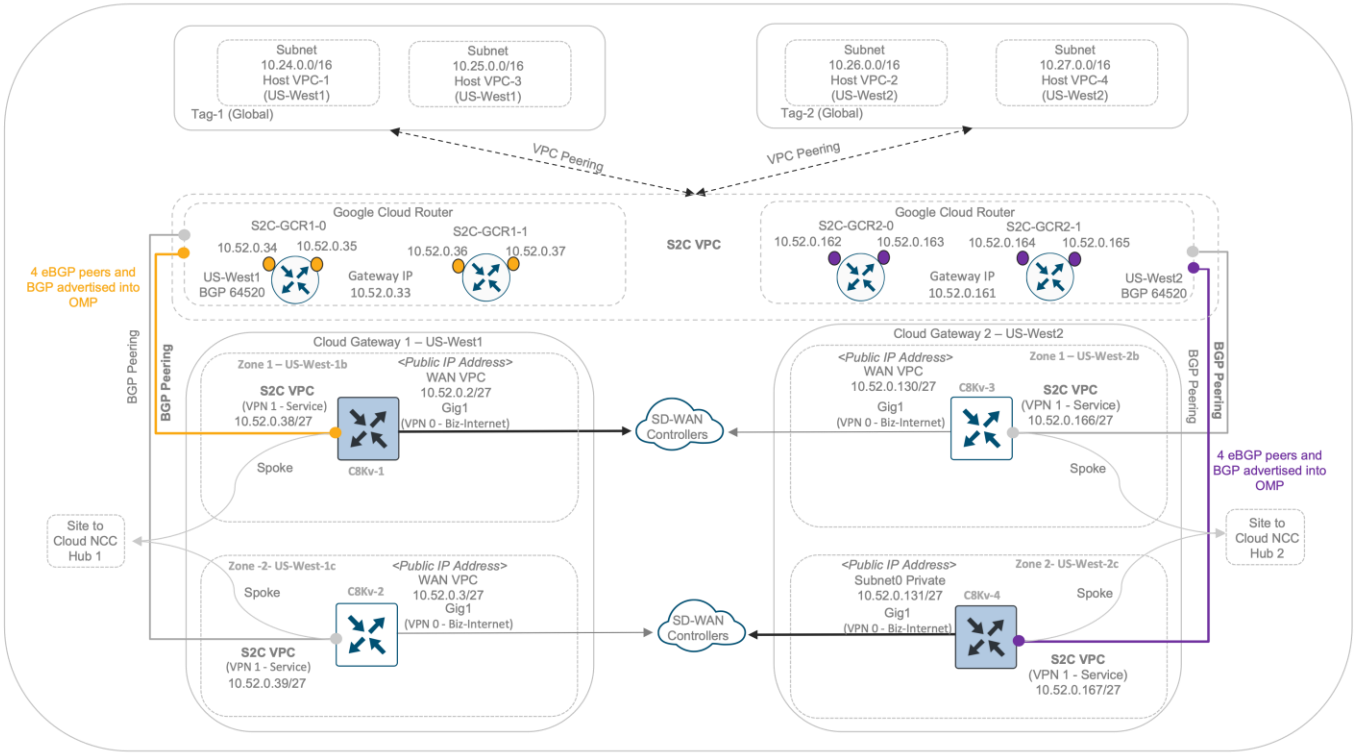

Figure 14. **BGP VPN 1 Routes**

```
CAT8KV-1#sh ip bgp vpnv4 vrf 1
BGP table version is 124, local router ID is 10.52.0.70
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf 1)
 *    10.24.0.0/16     10.52.0.33           10300             0 64520 ?
 *                     10.52.0.33           10300             0 64520 ?
 *                     10.52.0.33             100             0 64520 ?
 *>                    10.52.0.33             100             0 64520 ?
 *    10.25.0.0/16     10.52.0.33           10300             0 64520 ?
 *>                    10.52.0.33           10300             0 64520 ?
 *    10.26.0.0/16     10.52.0.33           10300             0 64520 ?
 *                     10.52.0.33           10300             0 64520 ?
 *                     10.52.0.33             100             0 64520 ?
 *>                    10.52.0.33             100             0 64520 ?
 *    10.27.0.0/16     10.52.0.33           10300             0 64520 ?
 *>                    10.52.0.33           10300             0 64520 ?
 *>   10.52.0.160/27   10.254.91.91          1000     50       0 ?
```

HOST VPC SUBNETS

GCR SUBNET ADVERTISED BY CAT8KV-3

© 2020 Cisco and/or its affiliates. All rights reserved.

All the data plane traffic destined to the applications hosted in Google Cloud from an on-premises SD-WAN branch site, connects to the WAN VPC. This in turn connects to customer hosted Google cloud VPCs through the site to cloud transit VPC.

Note - The term Google cloud hosted VPCs in this guide means either a host VPC, workload or applications VPCs that the customer hosted within Google Cloud.

Site-to-Site VPC Workflow

On enabling the site-to-site communication, a Site-to-Site (S2S) VPC is created and within the S2S VPC two pairs of redundant GCRs are created. Within the Catalyst 8000v devices provisioned in the CGWs, a second WAN interface is automatically configured under Interface Gigabit Ethernet 3 with color set to Private1. The tunnel interface can be of type GRE or IPsec with one private IP address assigned to it. This interface receives its subnet and associated firewall details based on resources automatically configured in the S2S VPC.

In the following example figure, we assume that two CGWs are provisioned one in US-West 1 region and the other in US-West 2 region. Within each CGW, site-to-site communication is configured. Therefore, two Google Cloud Routers (GCRs) are deployed in US-West 1 and two Google Cloud Routers (GCRs) are deployed in US-West 2, all a part of the same Site-to-Site VPC (S2S VPC). Within a pair of GCRs, one GCR acts as the primary virtual router and the other acts as the backup. Each Site-to-Site GCR (S2S GCR) is assigned with a BGP ASN, i.e., the entered BGP ASN Offset + 1.

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 15. **GCRs in S2S VPC**

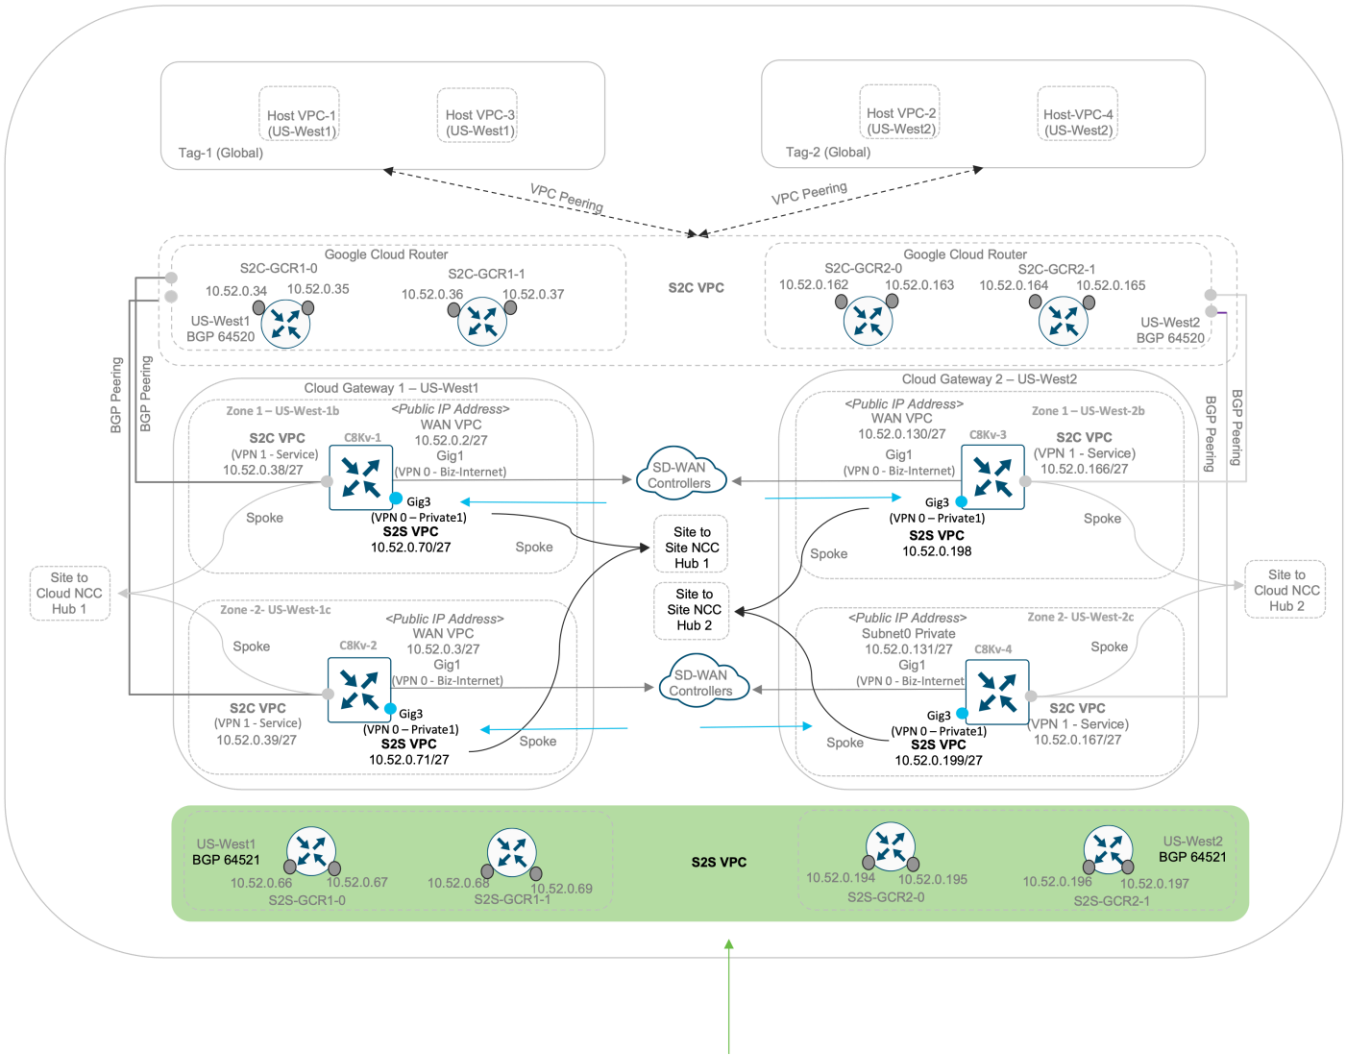

The GCR is the underlying routing engine that acts as a BGP peer in the GCP domain. Therefore, the GCR in the Site-to-Site VPC is programmed for eBGP sessions with the C8kv devices, therefore each C8kv forms BGP neighbors with 4 GCR IPs. The routes advertised through the neighbor include the S2S VPC GCR subnets.

The example figure below explains the BGP peering between one Catalyst8000v device and GCRs. In this deployment, the BGP ASN offset is set to 64520. Therefore, the GCRs are assigned a BGP ASN of 64521 (BGP ASN offset + 1) and Catalyst8000v devices in both the CGWs are assigned BGP ASN of 64530 (BGP ASN offset + 10).

The Catalyst 8000v device learns the GCR subnet from its 4 GCR eBGP neighbors with the next hop set as the GCR gateway IP.

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 16. **BGP Peering in S2S VPC**

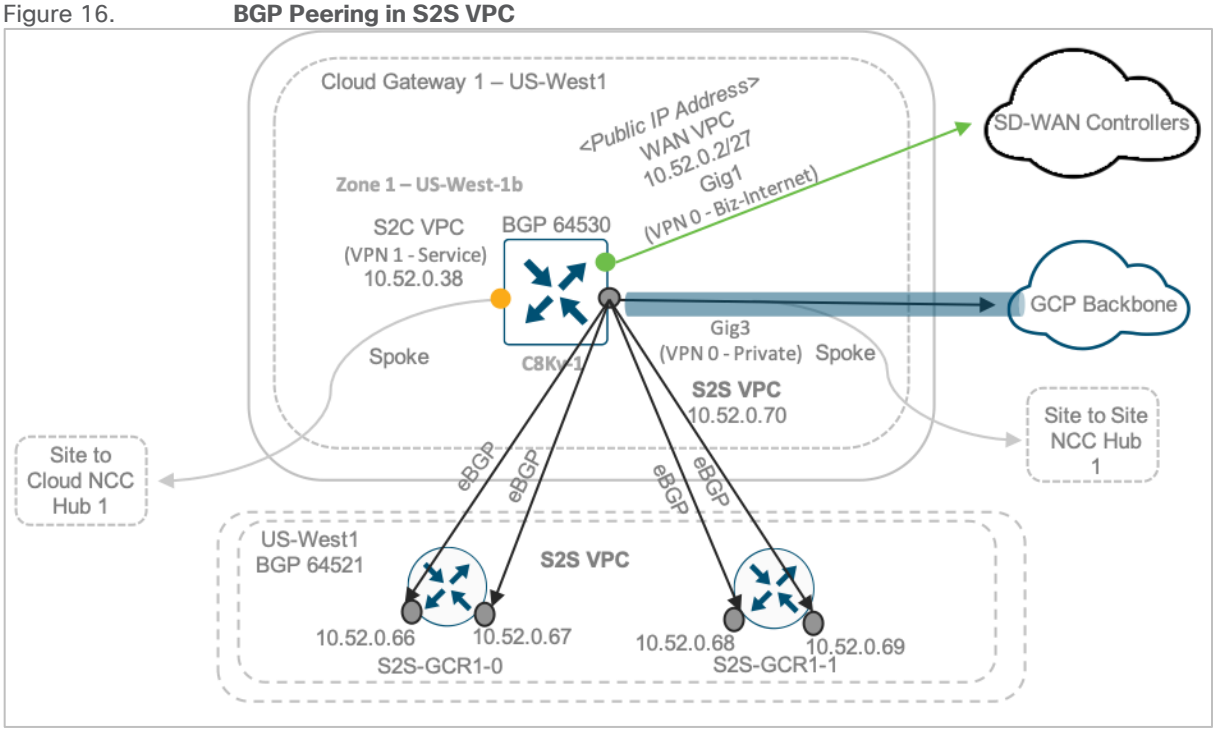

These neighbors advertise the entire GCR subnet to the Catalyst 8000v. The BGP routes learnt by the Catalyst8000v devices are advertised to the vSmart controller via OMP. The controller sends the advertised OMP routes to other SD-WAN Edge devices, to enable the WAN Edge devices located within on-premises or cloud hosted sites to establish tunnel connectivity with the Catalyst 8000v private link.

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 17. **GCR Subnet learnt in S2S Workflow**

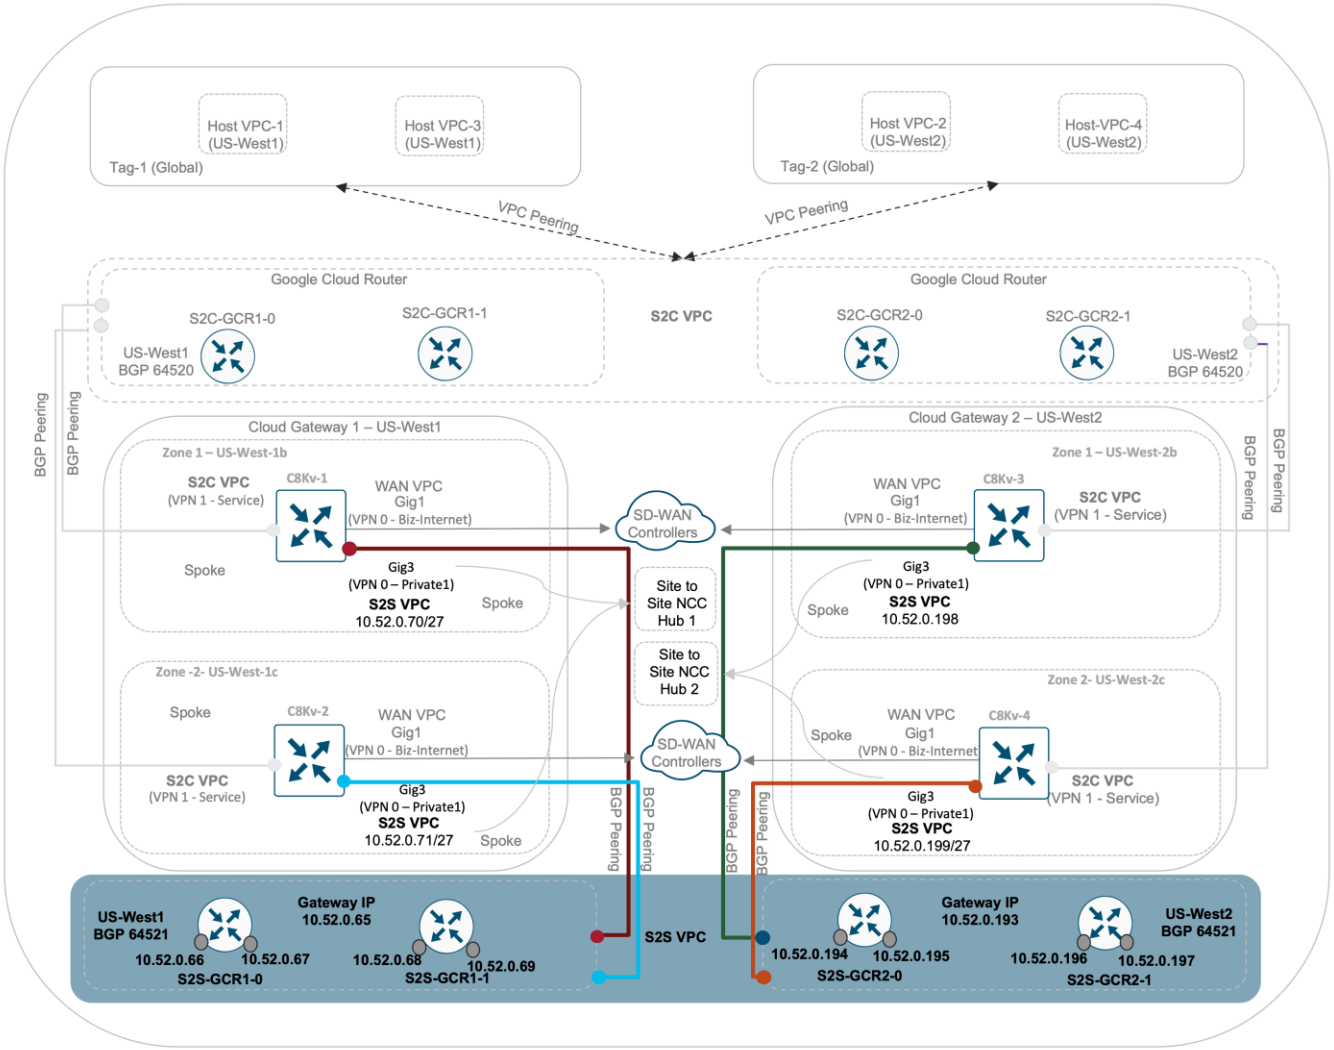

Within Google Cloud platform, the SD-WAN Cloud Gateway running Catalyst 8000v are modeled as Router Appliances in Google's domain and connect as spokes to the NCC Site to Site Hub. A Hub has a global scope while the SD-WAN CGW VMs are instantiated at regional level within the project. The Hub is responsible for inter-region routing using Google's backbone.

The Catalyst 8000v routers are created in different zones under the same region to meet the fault-tolerance needs. Between the two spokes the hub acts as a transit, therefore it allows for flow of the incoming data traffic through the Google Backbone.

Note: Given a scenario that your on-premises or cloud-hosted sites contain more than one WAN tunnel interface configured, say two WAN interfaces with two different colors assigned to it, MPLS and Private1, and the intent is to send critical traffic over the private link through your GCP global network and non-critical flows over other WAN tunnel links. For such design requirements, a centralized control policy can be configured to redirect traffic based on source IP/ port/ protocol to use the GCP backbone and the remaining non-critical traffic flows to exit via the remaining tunnels.

The following example illustrates a simple SD-WAN control policy, which redirects traffic to GCP using SD-WAN color Gold.

© 2020 Cisco and/or its affiliates. All rights reserved.

Figure 18. **Centralized Data Policy for S2S Workflow**

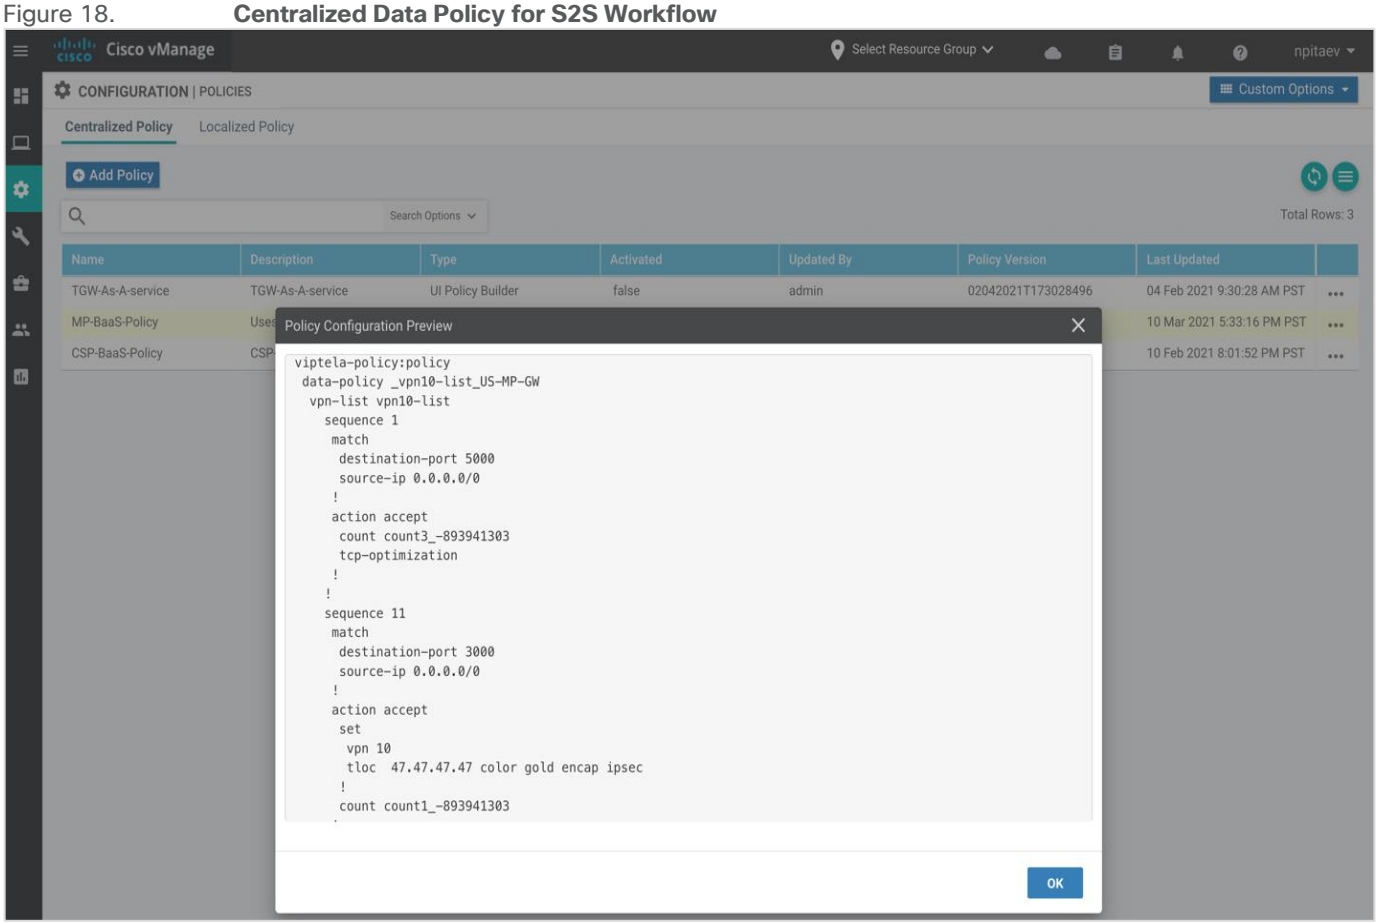

Please refer to the following Cisco Documentation for more details on SD-WAN Policies: https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/policies/vedge/policies-book/control-policies.html

Limitations and Best Practices

- The default GCP Catalyst 8000v template provisions one interface, that is the transport VPN 0 interface to connect to the public internet and the other two interfaces, that include second WAN transport interface for site-to-site traffic and service side interface for site to cloud traffic are configured at run time through autogenerated configuration.
- At least two free Cisco Catalyst 8000v virtual router UUIDs need to be available in Cisco vManage to bring up one Cloud Gateway.
- Only a maximum of two Catalyst 8000v devices can be deployed in a CGW, and each CGW must be deployed in different regions. The maximum number of CGWs provisioned in Google is based on the number of NCC regions available. Google Network Connectivity Center (NCC) is currently available in the following countries: https://cloud.google.com/network-connectivity/docs/network-connectivity-center/concepts/locations
- The workload VPCs should be set with a CIDR between 10.0.0.0 – 10.255.255.255 (10/8 prefix), 172.16.0.0 – 172.31.255.255 (172.16/12 prefix), and 192.168.0.0 – 192.168.255.255 (192.168/16 prefix)
- Validation checks are placed in the? solution, both at the cloud agent and NMS ports to make sure the CIDR agent is correct.

© 2020 Cisco and/or its affiliates. All rights reserved.

- Intra-VPC peering is enabled through tagging.
- Not all parallel actions are allowed since CGWs are global in nature. Anything after the first CGW is created, the rest of the CGW creations can be done in parallel. A second CGW cannot be created when the first one is being deployed or created. Similarly, deletion of all CGW's except the last one can be done in parallel.

Deploy - Cisco Cloud onRamp for Multi-Cloud using Google Cloud

This section covers the steps to deploy Cisco Cloud onRamp feature using Multi-Cloud.

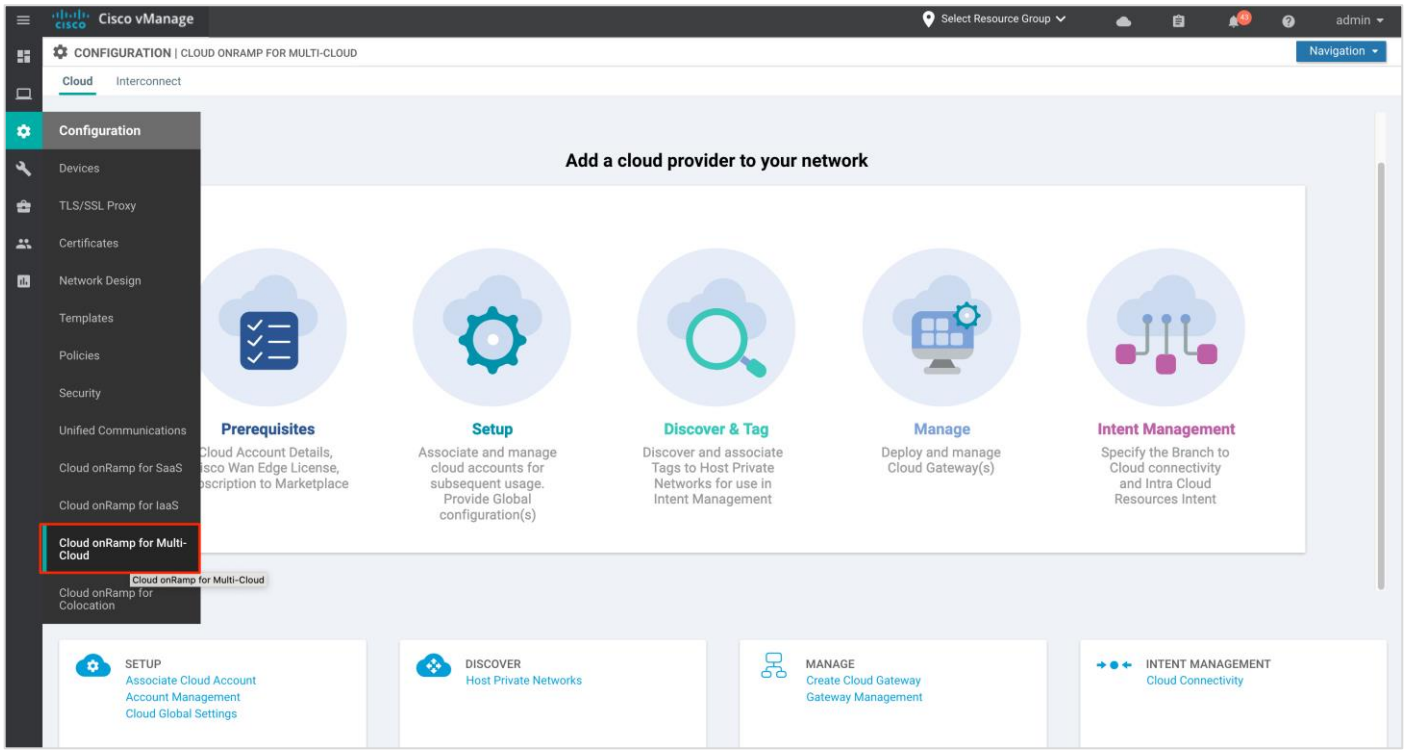

Configuration Workflows

- Prerequisites: This workflow includes all the prerequisites to enable and collect APIs, private keys etc. from the Google cloud account, followed by the steps to set up unused Catalyst8000v routers with device templates attached to it.
- Setup: This workflow includes the addition of the Google billing ID and the Google private key credentials within vManage, followed by the successful completion of the global settings required to successfully bring up Cloud Gateways.
- Discover &Tag: This workflow includes the discovery and association of Google Host VPCs to tags. You will use these tags later within the Intent Matrix page, where in you map tags to a service side VPN.
- Manage: This workflow creates and manages Cisco Catalyst 8000v virtual routers acting as cloud gateways. During this step, Cisco vManage will create three Google cloud VPCs, one for the WAN interface, one for the Site-to-Cloud (S2C) use case and the last one for the Site-to-Site (S2S) use case.
- Intent Management: The workflow maps the tags to the service side VPN. This allows on-premise branches to access Google cloud resources through the Site-to-Cloud VPC.

© 2020 Cisco and/or its affiliates. All rights reserved.

Process 1: Prerequisites Part 1 - Google Cloud Prerequisites

The procedures in this section list out the prerequisites to be completed within the Google cloud platform.

Procedure 1: Google Cloud Subscription

Ensure you have a Google Cloud account subscription as the account details are required to associate your Google cloud account with Cisco vManage.

Step 1. Ensure you have the following roles **Enabled** to your Google account.

- Service Account User
- Compute Instance Admin (v1)
- Compute Network Admin
- Compute Public IP Admin
- Compute Security Admin
- Hub & Spoke Admin
- Spoke Admin

Navigate to **IAM & Admin** > **Roles**, click on **+CREATE ROLE** to enable the required role.

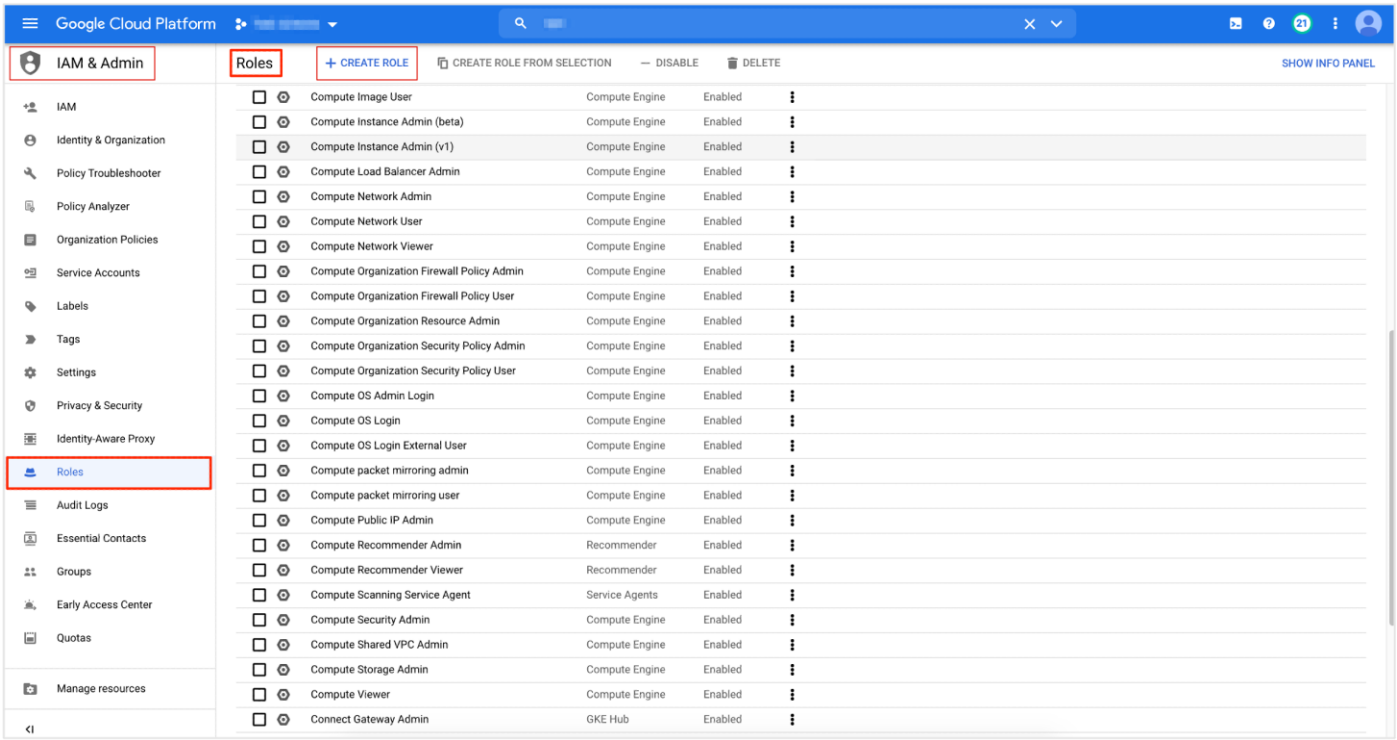

Procedure 2: Google Cloud API

Within your Google cloud account, the following APIs must be enabled, before you begin with the cloud onRamp workflow.

Ensure that following Google Cloud APIs are enabled in the relevant google project:

© 2020 Cisco and/or its affiliates. All rights reserved.

- Compute Engine API: This API is needed to create and run virtual machine, VPCs and firewall rules within the Google Cloud Platform.
- Cloud Billing API: This API is required to manage and validate billing for their Google Cloud Platform projects programmatically. In this guide, we have enabled both the cloud billing budget API and cloud billing API.
- Network Connectivity Center Alpha API: This API is needed to access, create and manage Network Connectivity Center hub and spoke resources.

Technical Tip
Your Google Cloud Project must have Network Connectivity Center Alpha APIs enabled. To enable this API, reach out to your Google contact.

Step 1. Navigate to **APIs & Services**. Under Dashboard, click on **+ENABLE APIS AND SERVICES**.

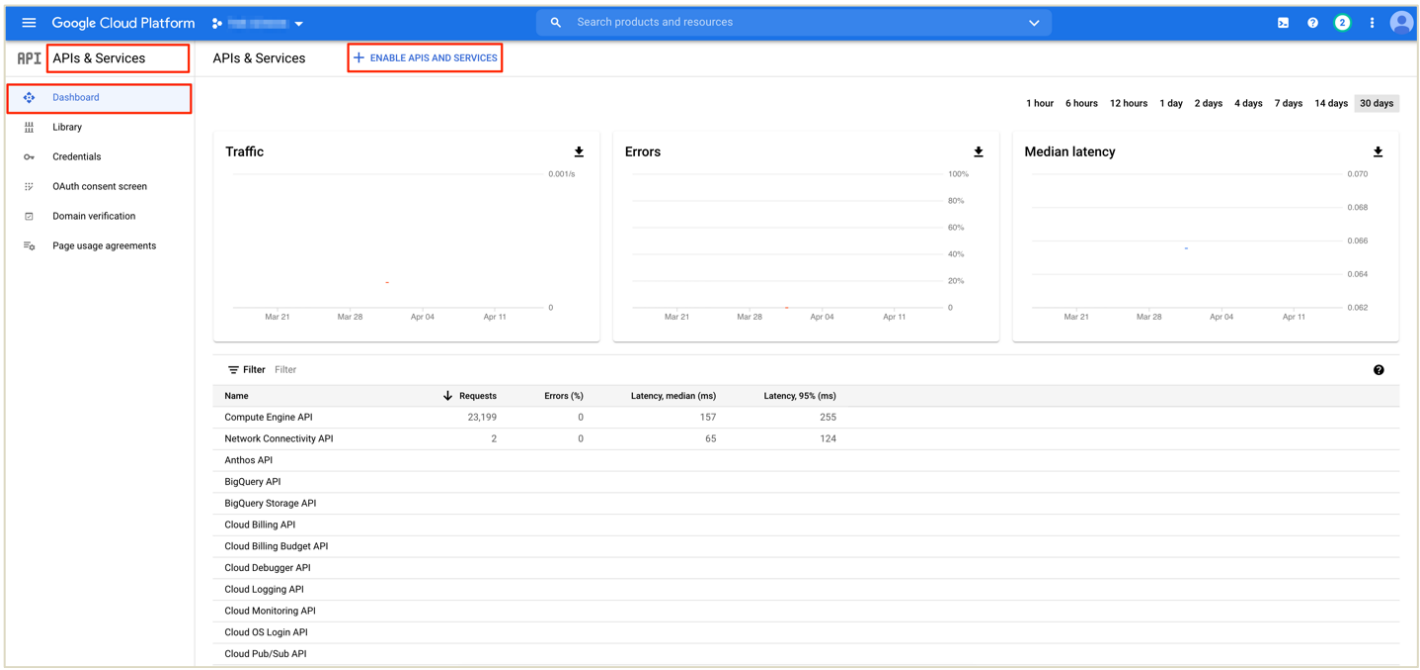

Step 2. Enter the APIs within the search bar and click **Enable**. The following screenshots display, the APIs enabled in this deployed.

© 2020 Cisco and/or its affiliates. All rights reserved.

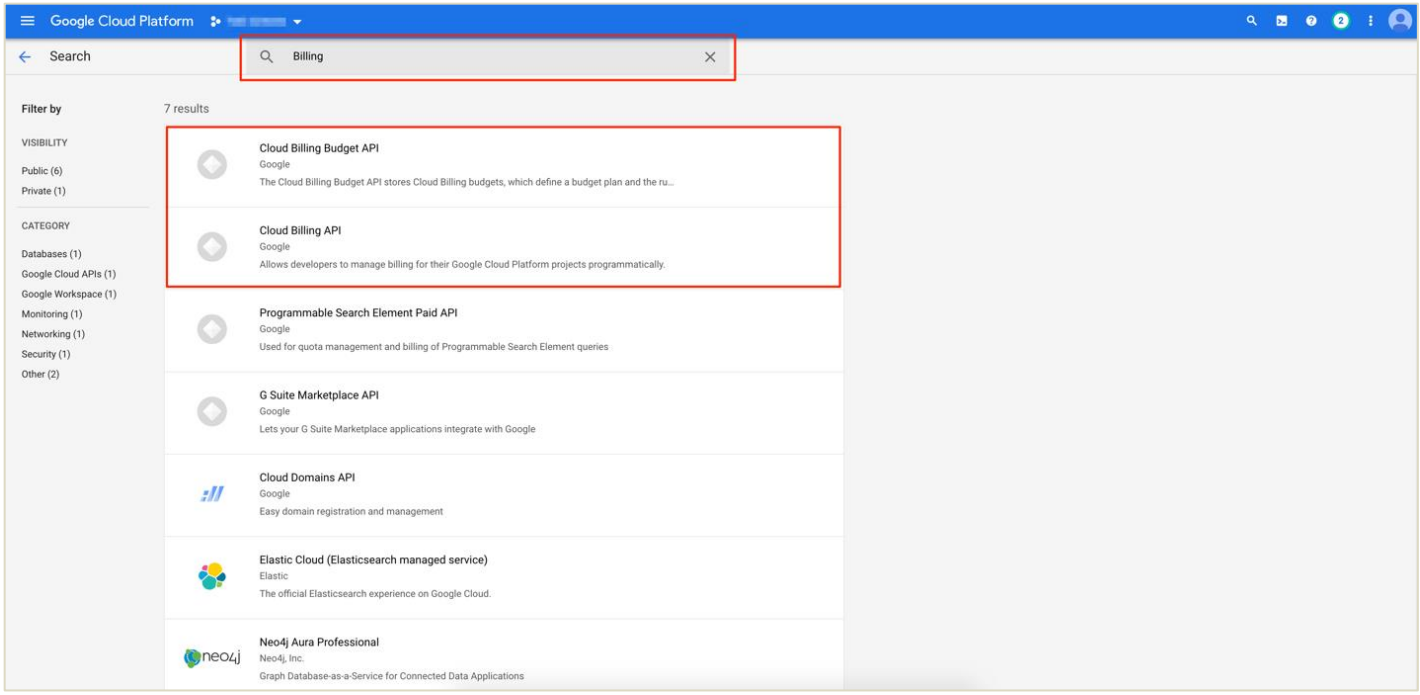

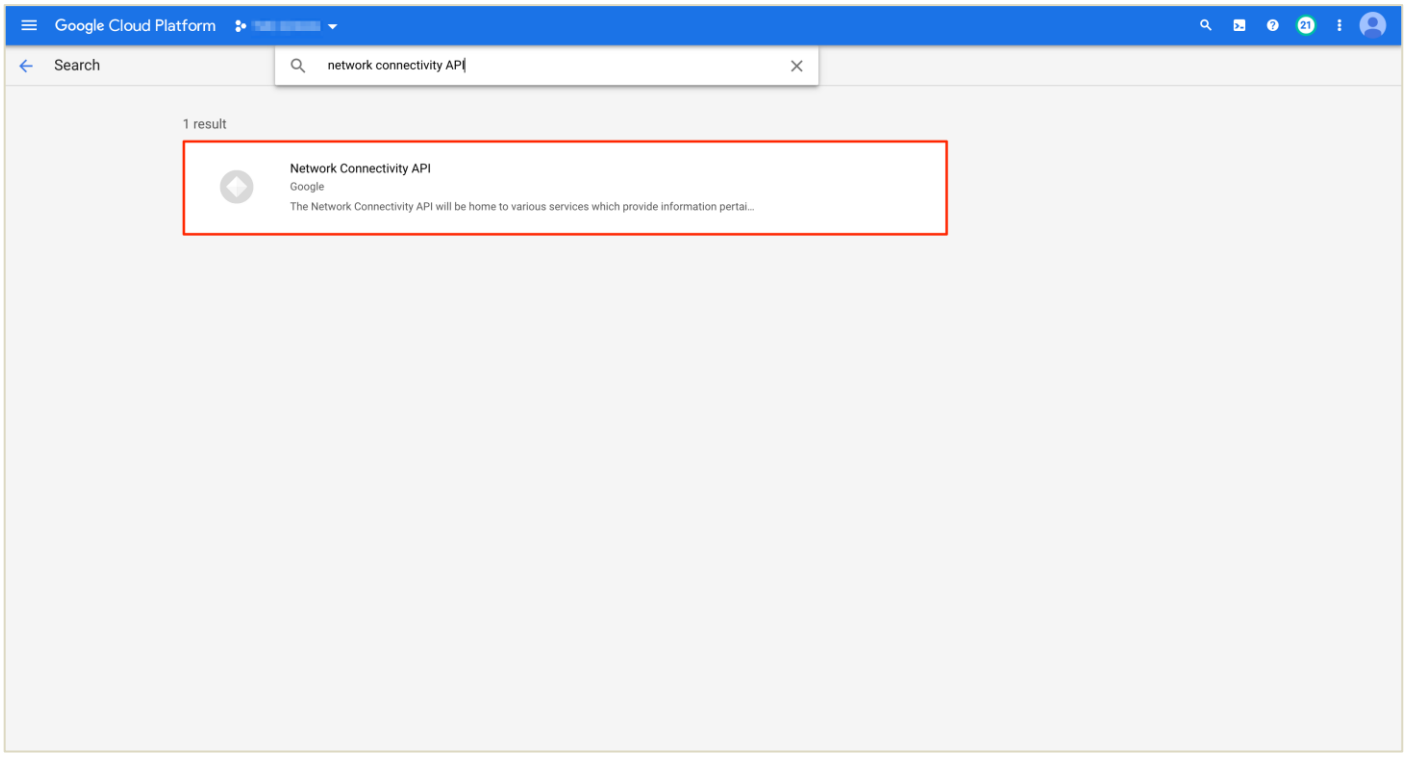

© 2020 Cisco and/or its affiliates. All rights reserved.

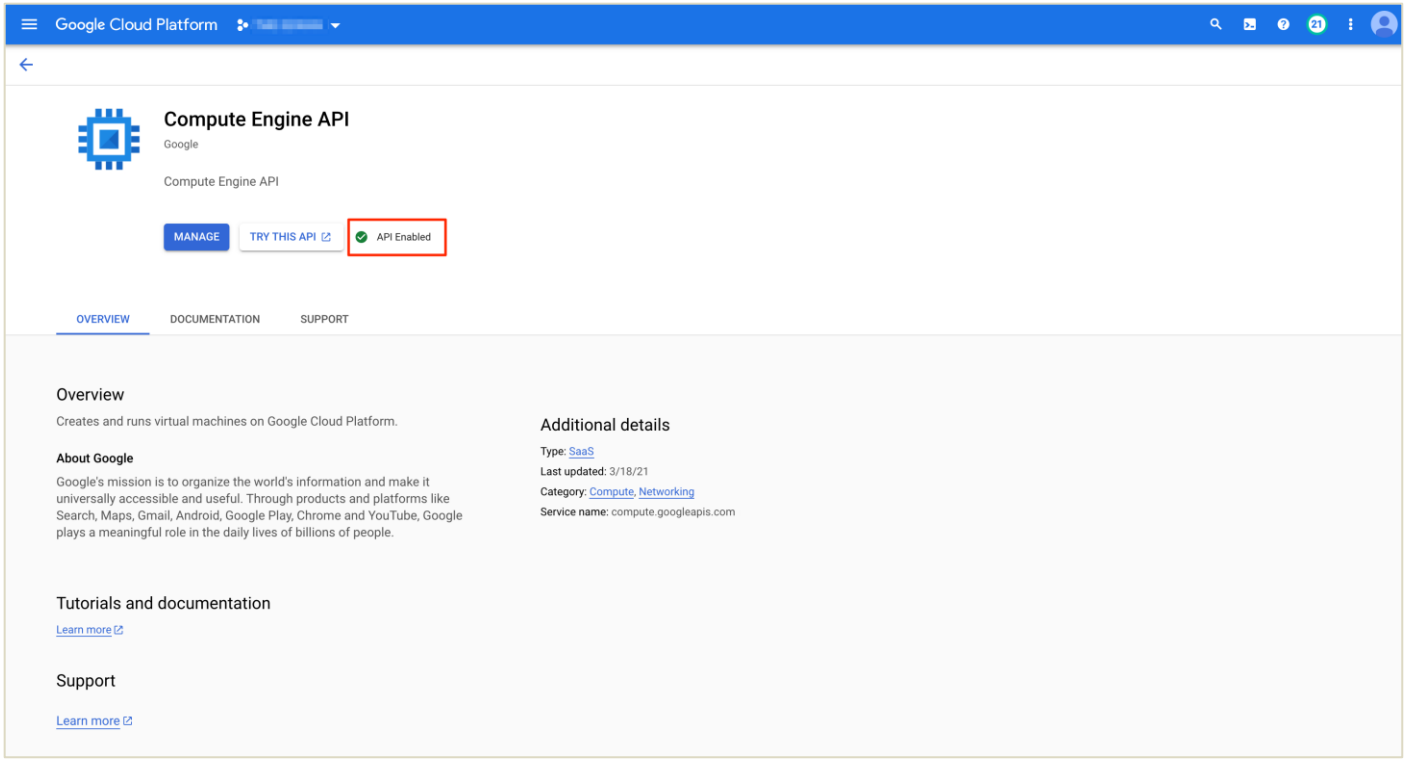

Step 3. While associating your Google Cloud account with your vManage, you can optionally add your Google Cloud billing ID. Navigate to **Google Cloud platform**, click on **Billing** from the hamburger ( ) drop-down menu and click on your **GO TO LINKED BILLING ACCOUNT**.

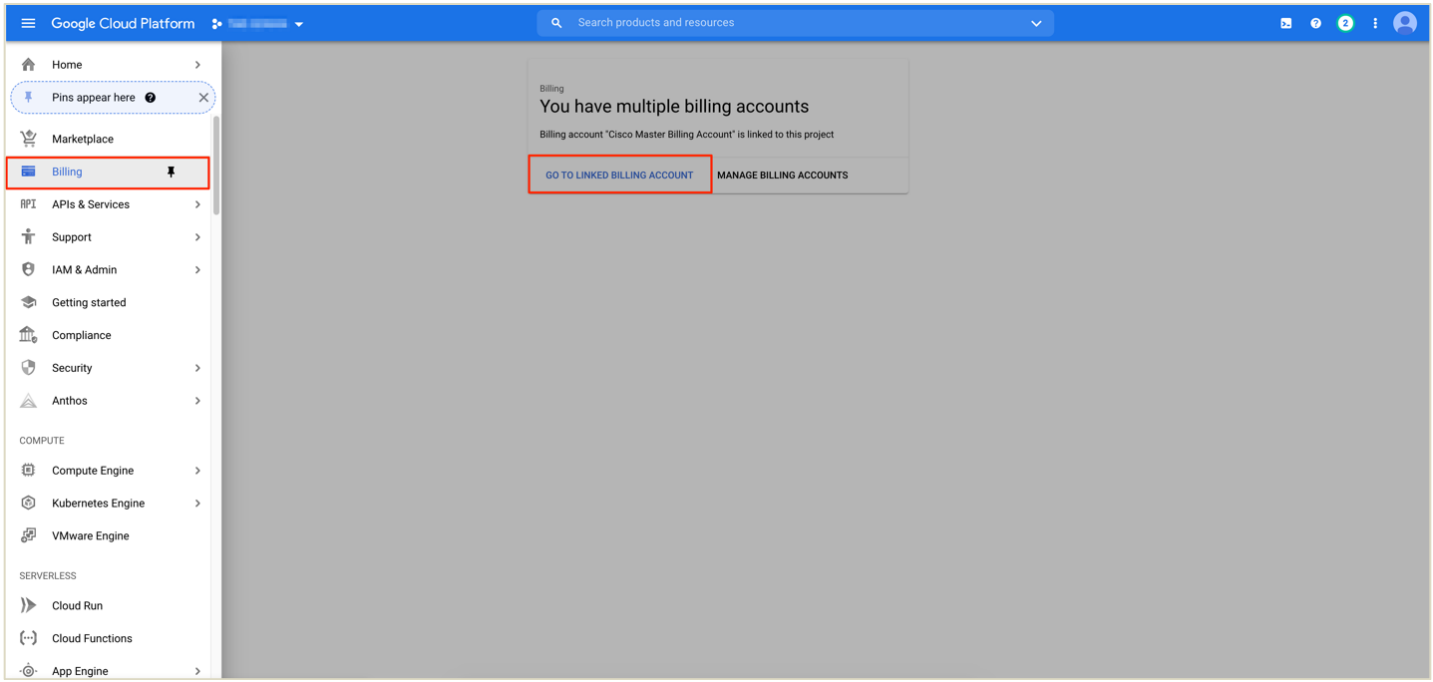

Step 4. Within **Billing**, select **Overview**. On the right side of the screen, you will find your billing account ID. Copy this ID as you need to enter this later while associating your Google Cloud account to your vManage controller.

© 2020 Cisco and/or its affiliates. All rights reserved.

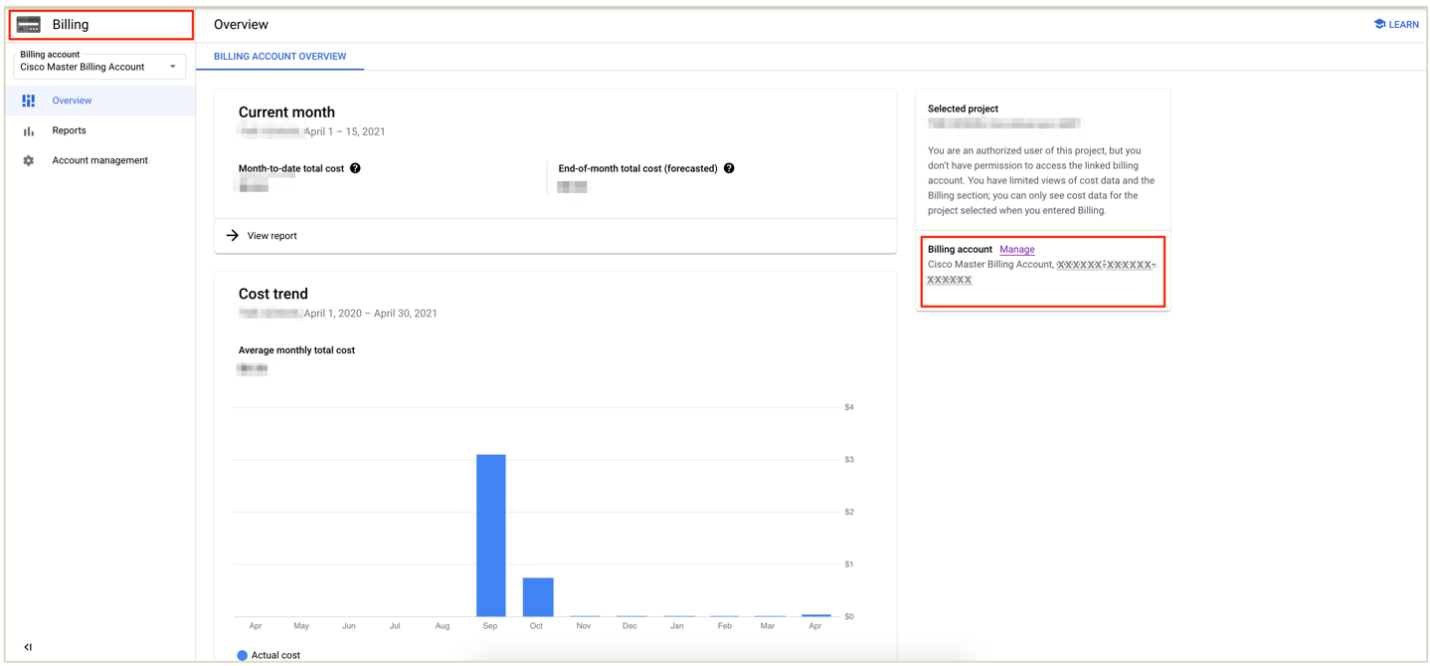

Procedure 3: Service Account and Private Key

In order to complete all the necessary cloud operations as a part of the cloud onRamp workflow, vManage NMS stores the Google account information in its database. To authenticate and accept the Google Account, vManage NMS uses the service account key method for authentication.

Therefore, if you do not have a service account in Google Cloud Platform defined for this purpose, you will need to create a new one and within the newly created service account a private key must be generated.

Technical Tip
Ensure that you create dedicated service accounts for each project that will have a host VPC.

Step 1. Navigate to **IAM & Admin** > **Service Accounts** and under **Keys**, click **ADD KEYS** and then select "Create new key" from the drop-down menu. to create and associate a new private key to your service account. The private key ID may be in **JSON** or **P12** - REST API formats. The format depends on the method of key generation.

The private key data returned is a base64-encoded string representation of the key. vManage will support both the aforesaid formats as input of the service account key. Once the type is chosen, click **Create**.

© 2020 Cisco and/or its affiliates. All rights reserved.

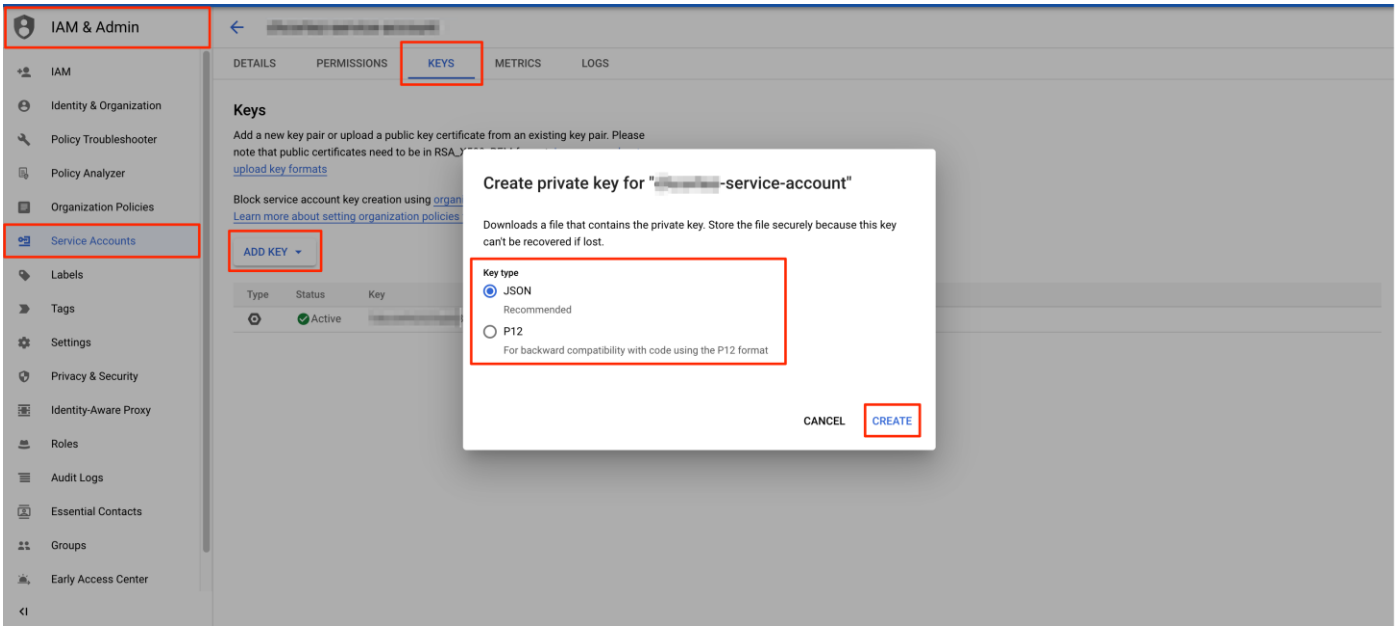

Step 2. Download or **Save** your key. This private key will be uploaded to your vManage while associating your Google account in the configuration workflow.

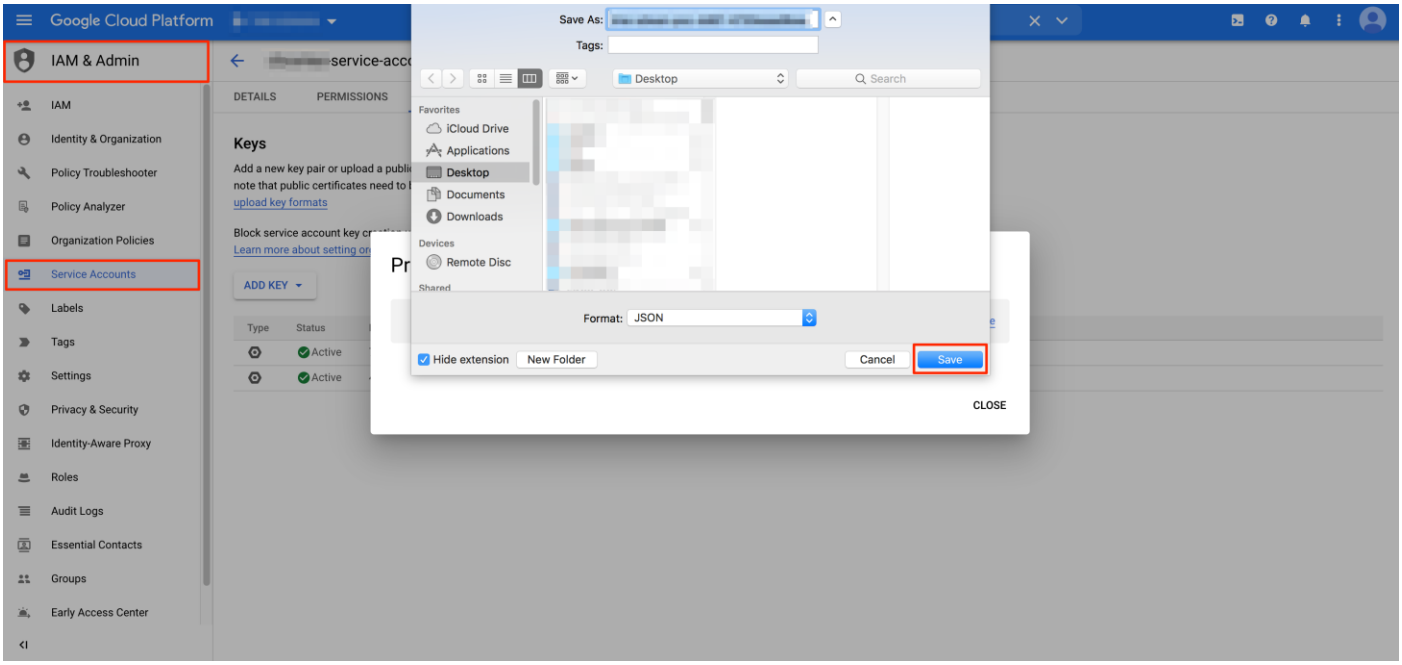

Note, once the key file has been downloaded, the same can be used on multiple vManage/domain controllers without repeating the authorization process for each controller.

Technical Tip
vManage will not allow updates to the credentials data. In order to update, user will have to first deregister the account with vManage and register afresh.

© 2020 Cisco and/or its affiliates. All rights reserved.

Procedure 4: Firewall Requirements

The following TCP, UDP and ICMP ports are automatically allowed for incoming traffic in Google Cloud upon completion of site-to-cloud and site-to-site workflows and each of the firewall rules are attached to the Site-to-Cloud/ Site-to-Site/ WAN VPCs. In Google Cloud Platform (GCP), all the ports for egress or outgoing traffic are automatically opened. The following example lists out the ingress **Protocols/Ports**, allowed within the Google Cloud firewall rules.

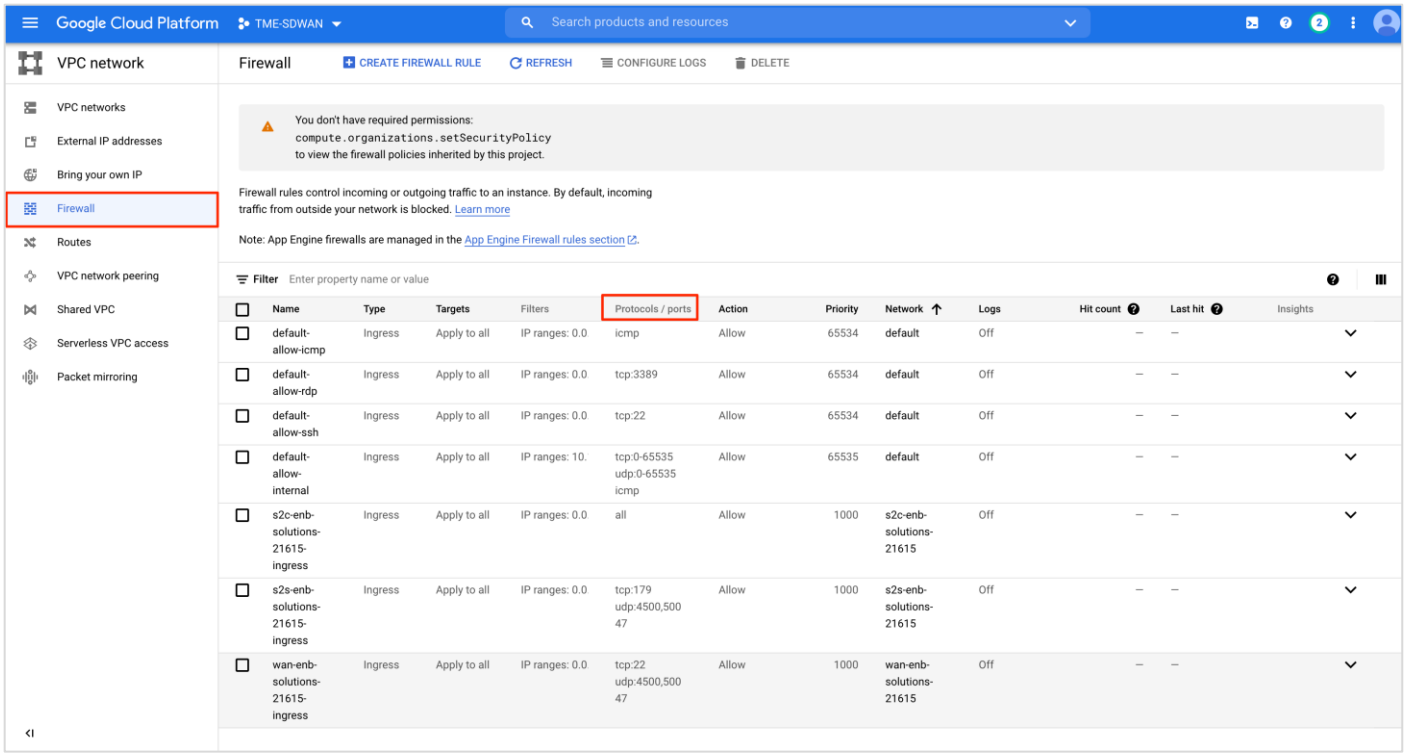

The protocols allowed automatically on WAN VPC,

- TCP: 22
- UDP: 4500, 500

The protocols allowed automatically on S2C VPC,

- ALL

The protocols allowed automatically on S2S VPC,

- TCP: 179
- UDP: 4500, 500

And by default, for all VPCs these ports are allowed,

- TCP: 3389, 22, 0-65535
- UDP: 0-65535
- ICMP

The DTLS control connections to vBond, vSmart and vManage are initiated from the WAN Edge device. Once the control connections are initiated from the WAN Edge device (Catalyst8000v), Google cloud maintains a stateful session to the destination IP/ port/ protocol to complete the authentication process and bring up of the WAN Edge devices. Therefore, no ports need to be manually opened by the user in the Google Firewall to establish control and data plane connections.

© 2020 Cisco and/or its affiliates. All rights reserved.

However, make sure to open DTLS/TLS ports within the on-premises firewall to establish control and data plane connections with your cloud devices. For details on the ports to be opened, refer to the Cisco SD-WAN Design Guide.

Process 1: Prerequisites - Part 2 - vManage Prerequisites

This section focuses on all the prerequisites within the vManage NMS.

Procedure 1: vManage Internet Access

Step 1. Make sure that your vManage server has access to the Internet and that it has a DNS server configured so that it can reach Google Cloud. To enable the DNS server configuration, you can configure this either within the VPN feature template associated with your vManage device template or via CLI.

Option 1: If your vManage is configured using the vManage device templates, then a DNS server configuration is added in the vManage VPN 0 feature template. To configure a DNS server in vManage VPN 0, enter the IP address of a DNS server and then save the edited feature template.

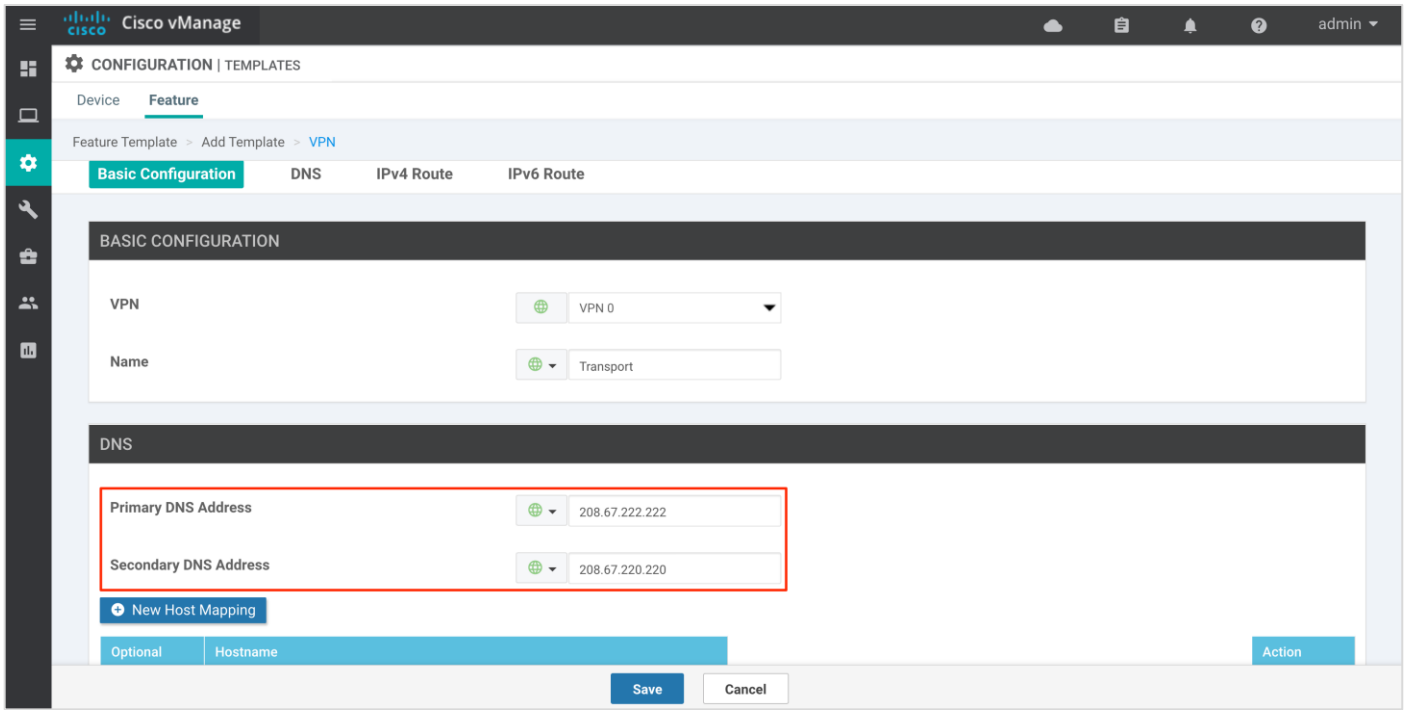

Option 2: If your vManage is configured manually via CLI, then log into the vManage GUI and navigate to **Tools** > **SSH Terminal**. Click on the vManage server from the device group, log in with the admin username and password, and enter the following commands:

```
vManage(config)# vpn 0
vManage(config-vpn-0)# dns 208.67.222.222 primary
vManage(config-vpn-0)# dns 208.67.220.220 secondary
vManage(config-vpn-0)# commit
```

Procedure 2: vManage Time Server Synchronization

Step 1. Ensure that the vManage NMS server is synchronized to the current time. To check the current time, click the Help (?) icon in the top bar of the vManage screen or issue a “show clock” via CLI.

© 2020 Cisco and/or its affiliates. All rights reserved.

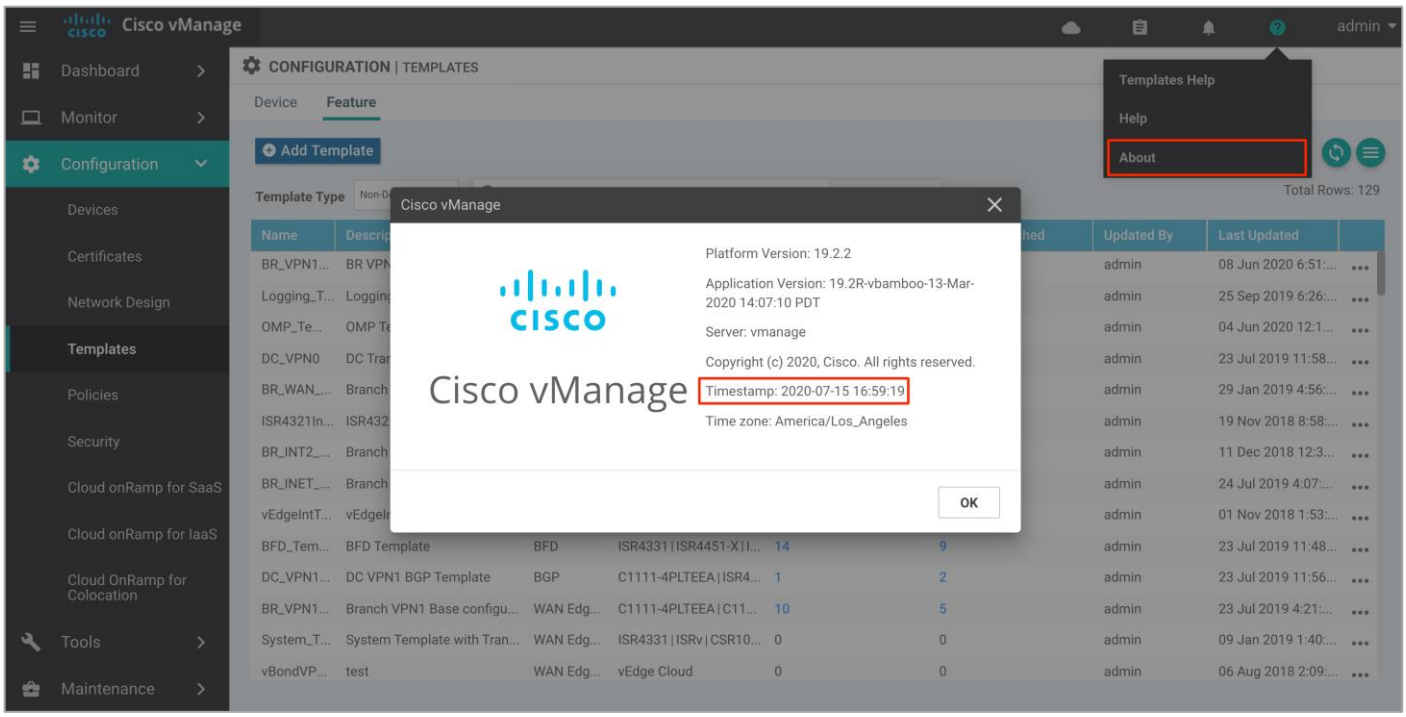

Step 2. The Timestamp field shows the current time. If the time mentioned is incorrect, configure the vManage server's time to point to an NTP time server, such as the Google NTP server. To do this, either an NTP feature template can be associated within the vManage device template, or you can configure it manually via CLI.

Option 1: To configure or update the vManage NTP feature template, enter the **Hostname/ IP Address** of an NTP server, and then attach the new or updated feature template within the vManage device template.

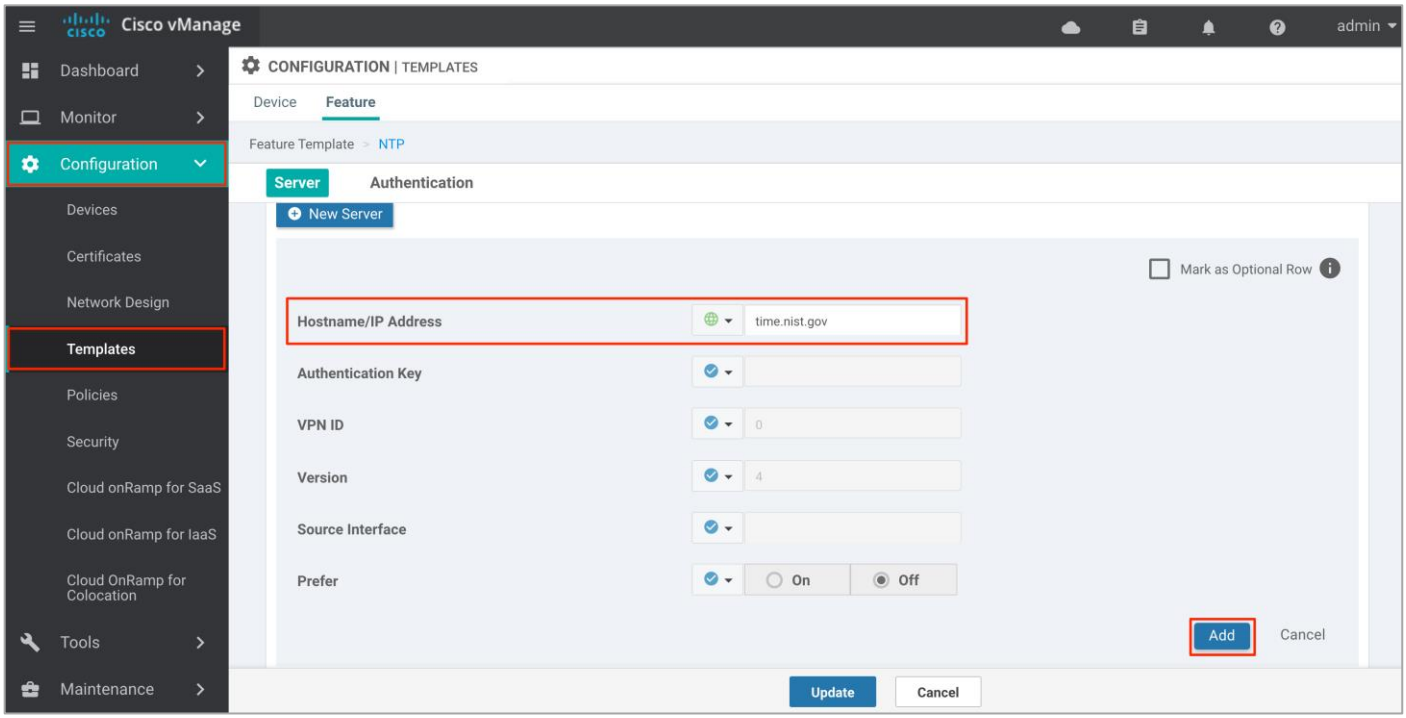

Option 2: To configure an NTP server via CLI, login to the vManage GUI and navigate to **Tools** > **SSH Terminal**. Click on the vManage server from the device group, log in with the admin username and password, and enter the following commands:

```
vManage# config t
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
vManage(config)# system
vManage(config-system)# ntp
vManage(config-ntp)# server time.nist.gov
vManage(config-server-time.nist.gov)# version 4
vManage(config-server-time.nist.gov)# commit
Commit complete.
```

Procedure 2: Verify you have at least four unused Cisco Catalyst 8000v routers in vManage

To complete the site-to-cloud and site-to-site use case, four Catalyst 8000v devices are deployed in Google cloud. Two of the routers are deployed in region 1 and a second pair of devices are deployed in a different region, Region 2. Each of the WAN Edge devices are placed in separate zones (automatically?) for fault-tolerance with high availability.

In this workflow, a pair of unused Catalyst8000v devices are deployed in US-West1 region and a second pair of unused Catalyst8000v devices are deployed in US-West2 region. Therefore, a total of four unused Catalyst8000v devices must be attached to a vManage device template.

Begin by adding Catalyst 8000v routers within your vManage device list.

Step 1. Login to Cisco **Plug and Play portal Connect** via URL - https://software.cisco.com/#pnp-devices . Under the **Devices** tab, click **+ Add Software Devices** to add the devices to the portal.

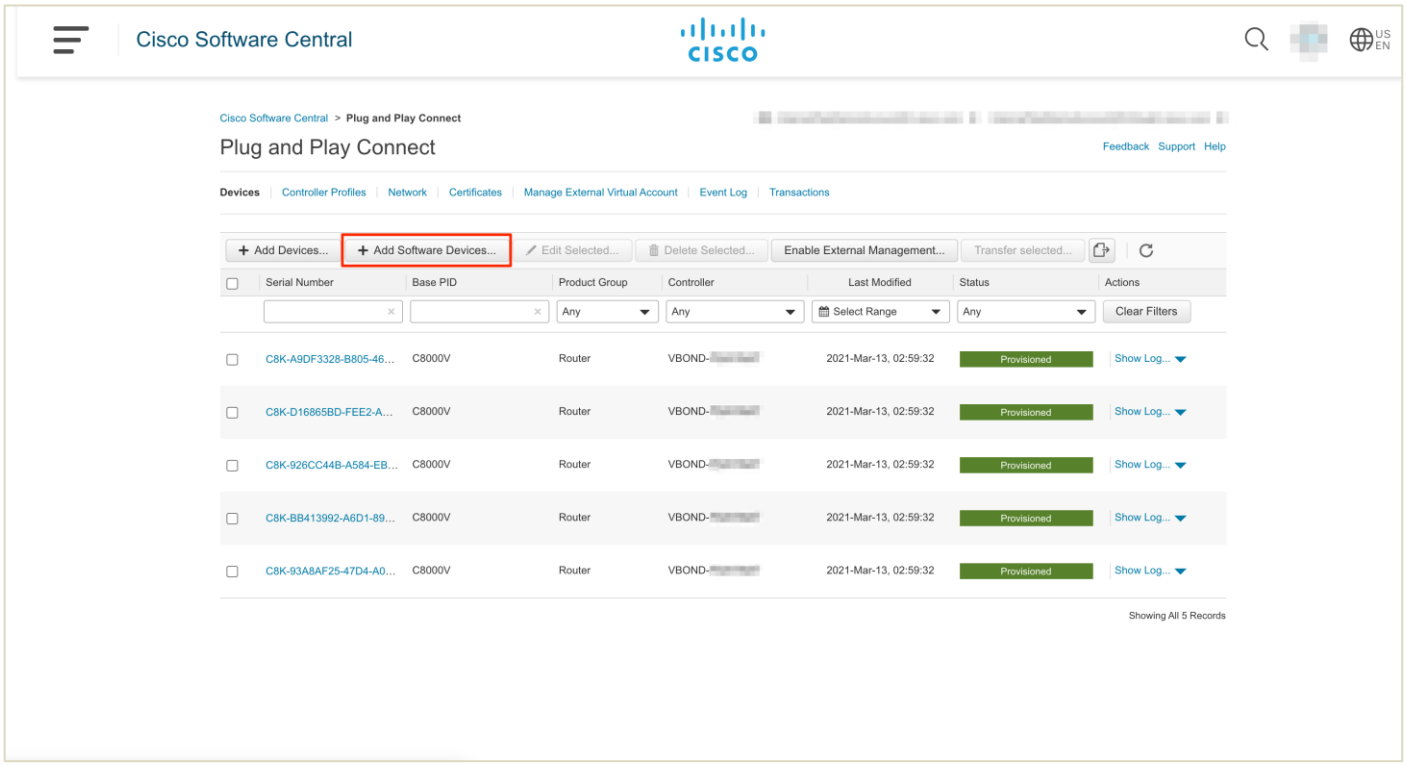

Step 2. To add Catalyst 8000v devices, enter **Base PID** as **CATALYST 8000V** and **Quantity** as 4 or more, and select your **Controller Profile** from the drop-down option. Click **Save**, then **Next.** then **Submit**. Click **Done**.

© 2020 Cisco and/or its affiliates. All rights reserved.

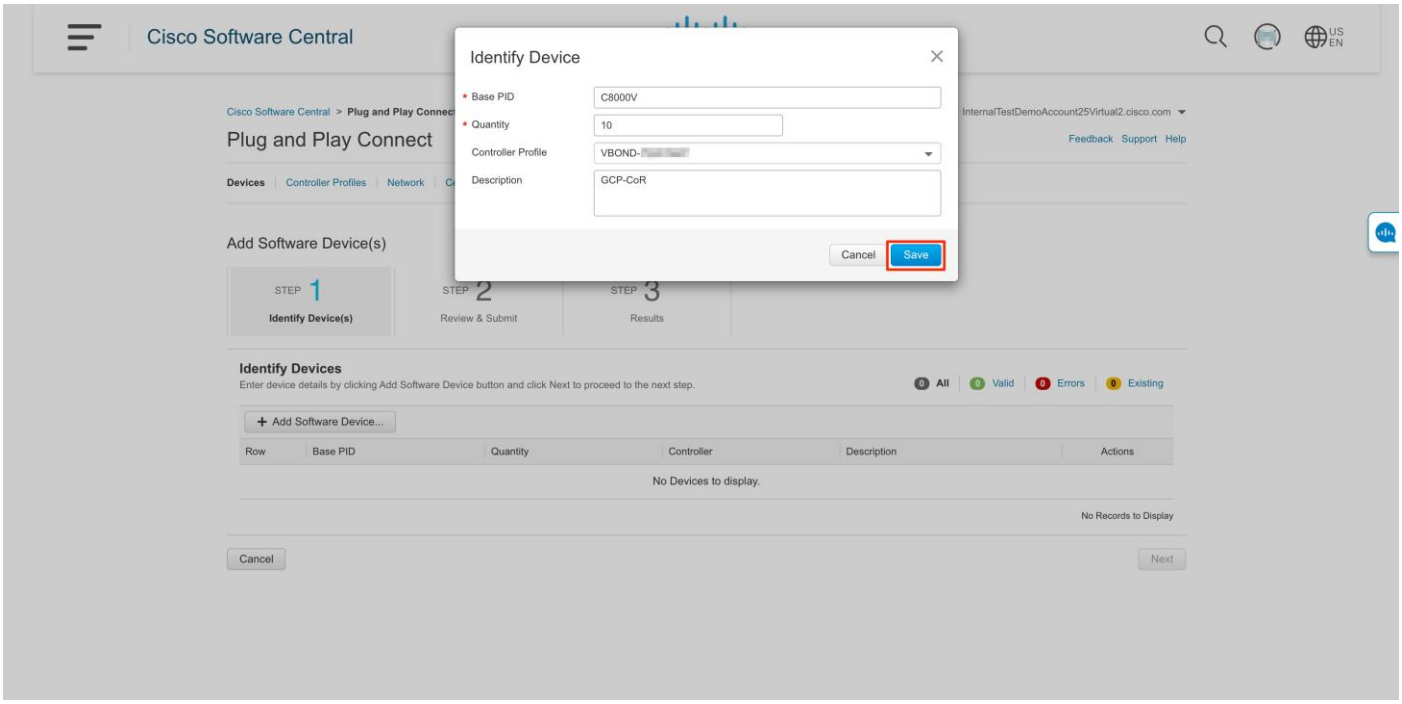

Step 3. Navigate to **Configuration** > **Devices** > **Sync Smart Account** to sync your vManage controllers to the **Cisco Plug and Play** portal. This automatically adds the newly generated devices to the vManage device list.

To add the new devices to your vManage controller manually, within the Plug and Play portal, go to the **Controller Profiles** tab, download the **Provisioning File** available within right side corner of the **Controller Profile** page. Once downloaded, navigate to vManage NMS server and go to **Configuration** > **Devices** > **Upload WAN Edge List**. Do not forget to click on the checkbox "**Validate the uploaded vEdge List and send to controllers**" before uploading the device list.

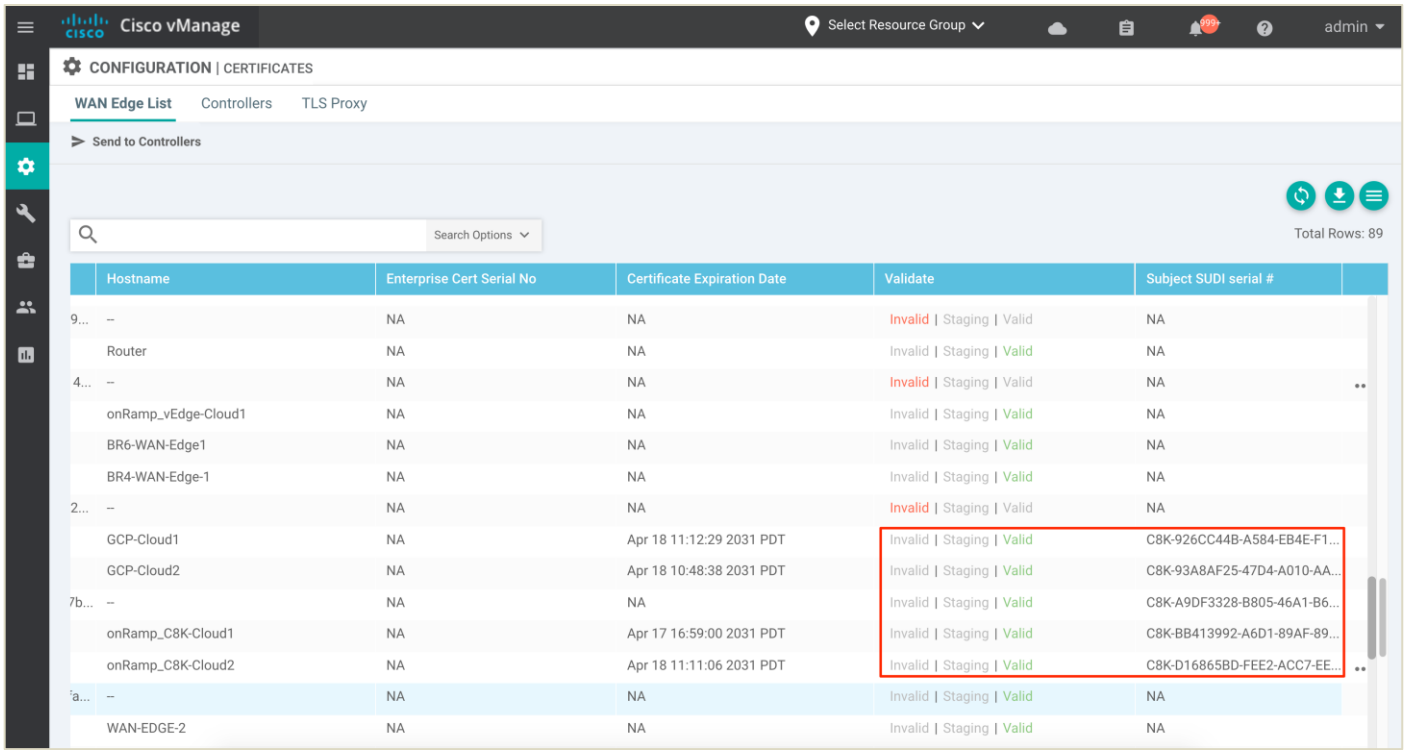

© 2020 Cisco and/or its affiliates. All rights reserved.

Step 4. Within the vManage dashboard, navigate to **Configuration** > **Templates** and select the **Device** tab.

Step 5. Find the desired GCP default template (**Default_GCP_CATALYST 8000V_Template_V01**).

An example is shown in the following figure.

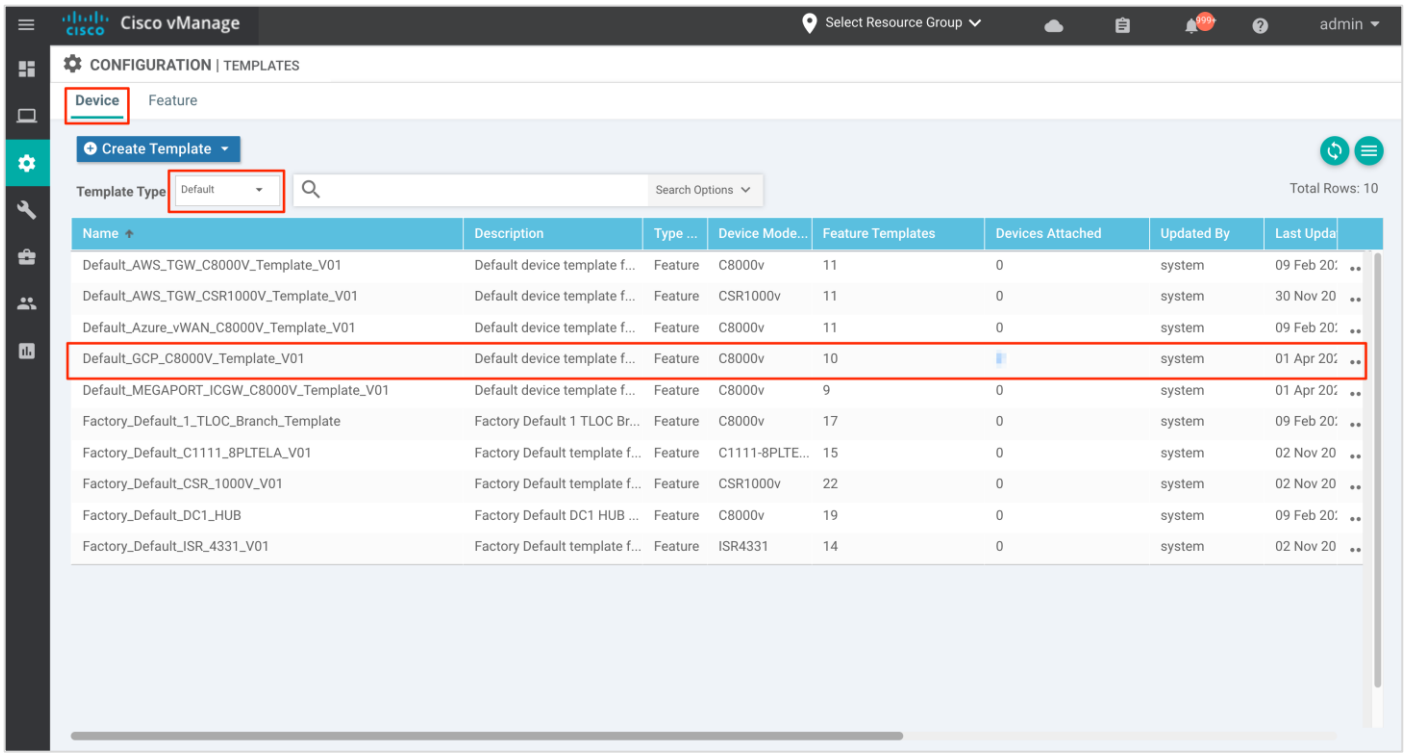

Step 6. Within this drop down, you can choose to either **copy** and rename the device template, and later attach devices under the non-default template you've generated or click on attach devices to add devices under the default template. The following screenshot explains the former option.

Click on the three dots **(...)** located on the right side and click **Copy**.

© 2020 Cisco and/or its affiliates. All rights reserved.

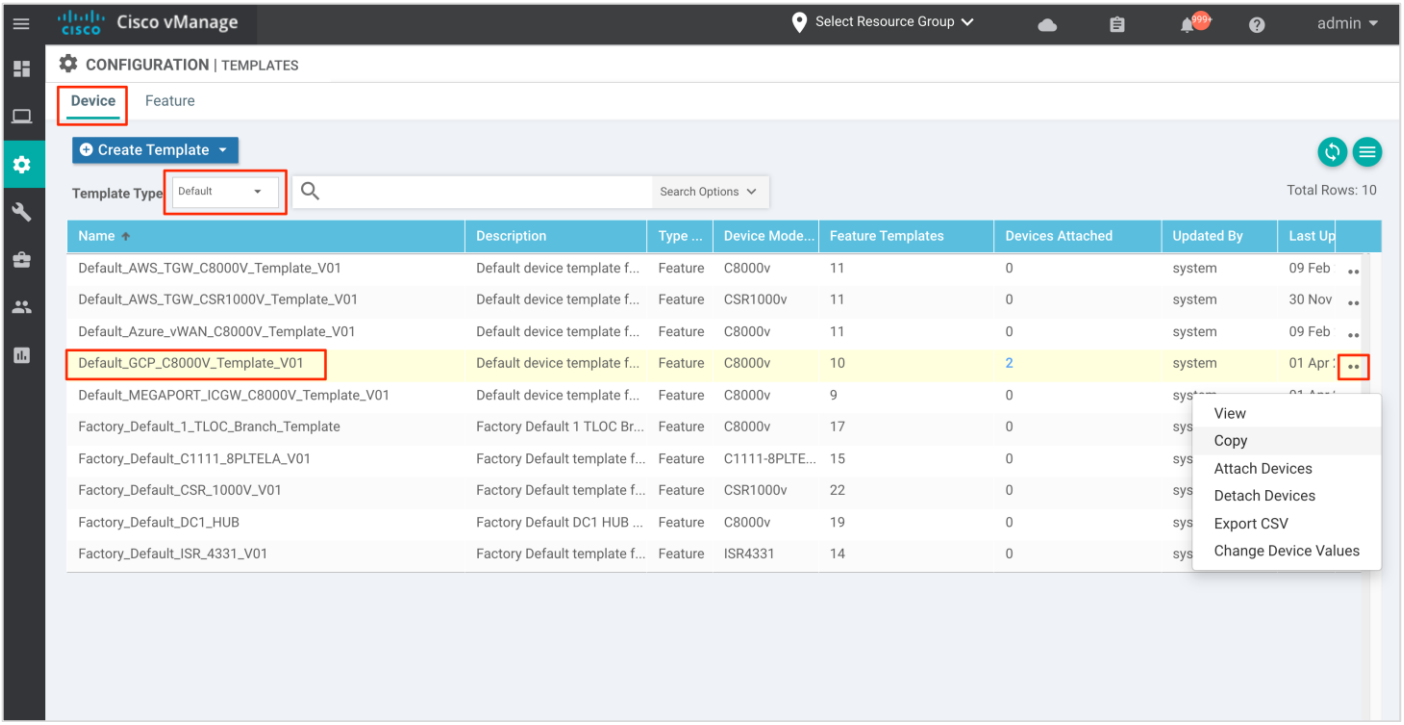

Step 7. The template is renamed to GCP_CATALYST 8000V_Template_V01 and click **Copy** again.

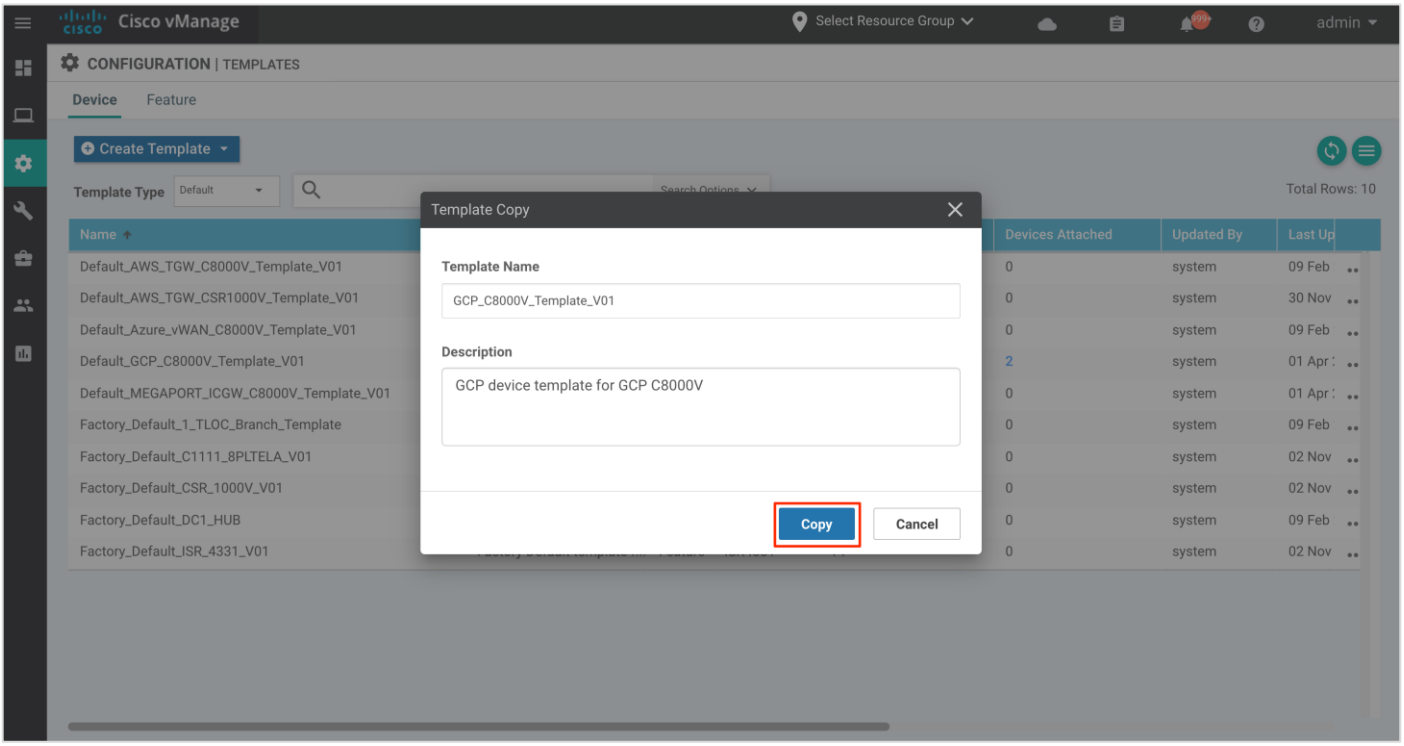

Step 8. Under the **Device** section, click on the **Template Type** drop-down list and select **Non-Default**. Click on the three dots (**...**) next to the non-default template, **GCP_8000V-Template_V01** and click **Attach Devices**.

© 2020 Cisco and/or its affiliates. All rights reserved.

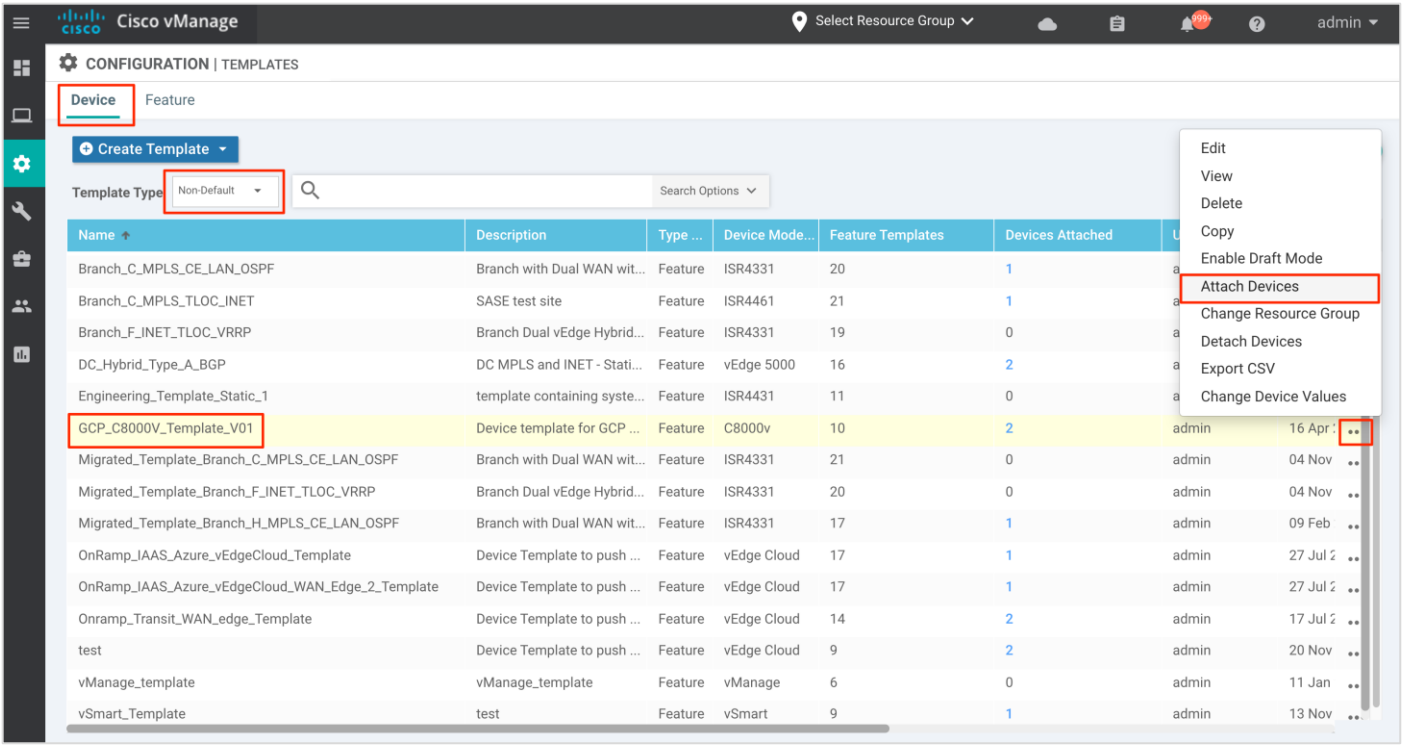

Step 9. A pop-up window listing the available devices to be attached to this configuration will appear. The list of available devices contains either the hostname and IP address of a device if it is known through vManage; or contains the chassis serial number of the devices that have not yet come up on the network and are unknown by vManage. Cisco Catalyst8000v routers are assigned a chassis serial number although there is no physical chassis.

Select the devices you want to apply the configuration template to and select the arrow to move the device from the **Available Devices** box to the **Selected Devices** box.

You can select multiple devices at one time by simply clicking each desired device.

© 2020 Cisco and/or its affiliates. All rights reserved.

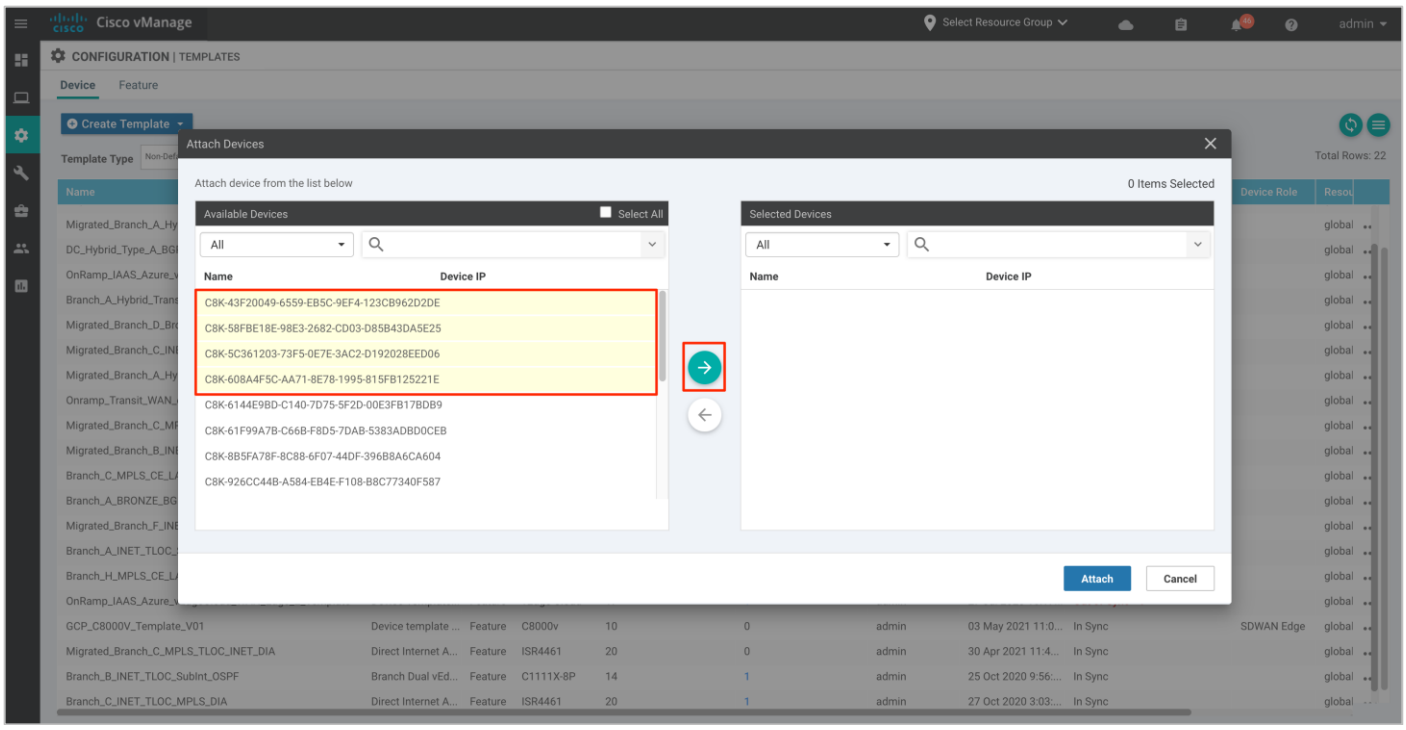

Step 10. Click on the **Attach** button.

A new screen will appear, listing the devices that you have selected.

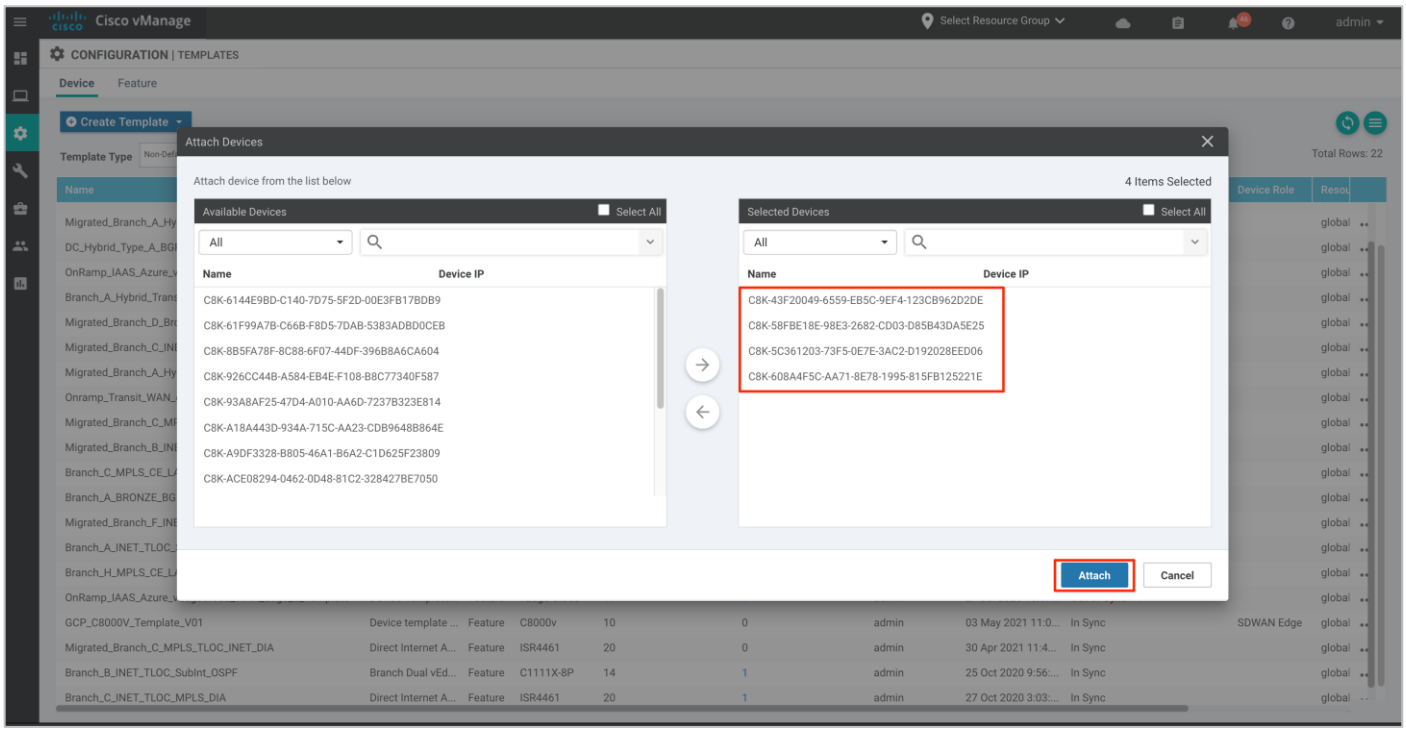

Step 11. Click on the three dots (**...**) located to the far right of each device template, and from the drop-down menu select **Edit Device Template** or simply enter the values directly into each of the column.

An example is shown in the following figure.

© 2020 Cisco and/or its affiliates. All rights reserved.

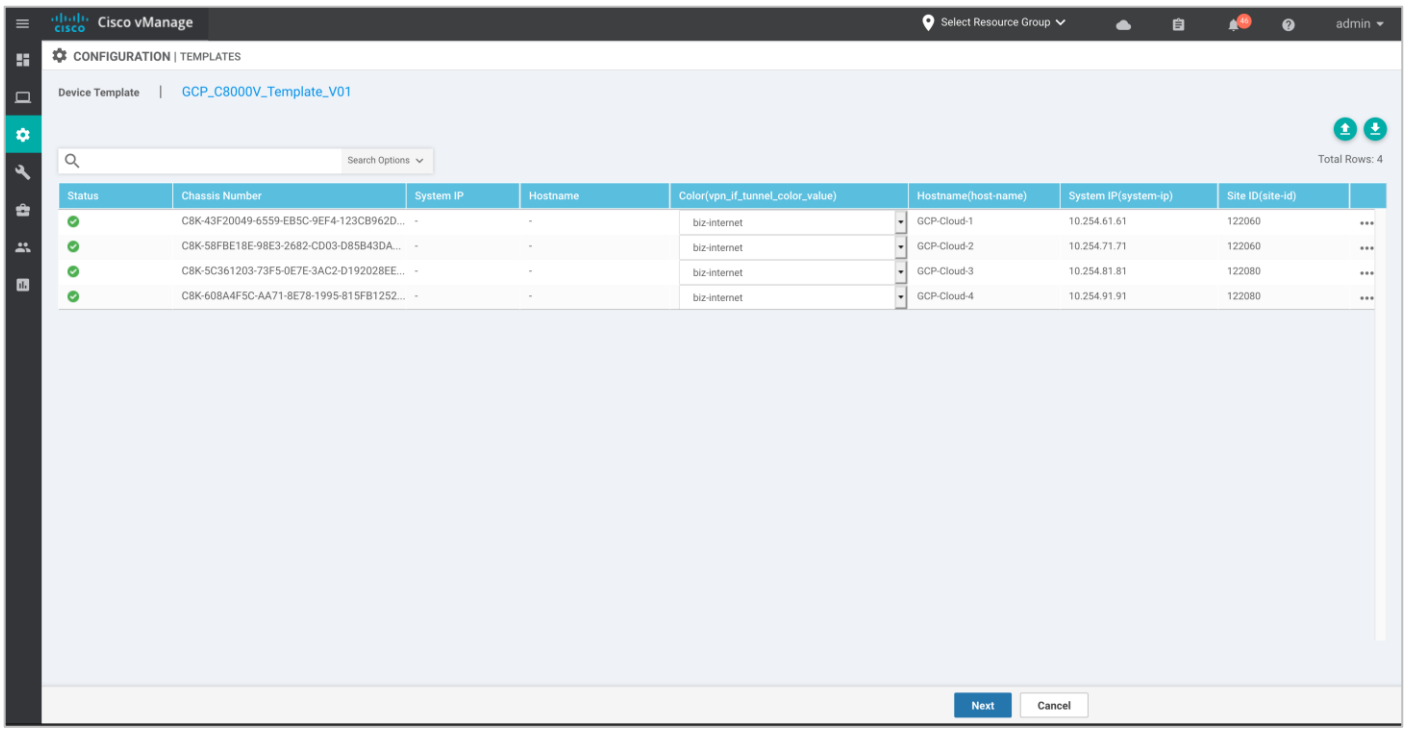

Step 12. Fill in the values of the variables in the text boxes.

The following template is deployed on Chassis Number – **C8K-43F20049-6559-EB5C-9EF4-123CB9262D**

Variables	Value
Color (vpn_if_tunnel_color_value)	Biz-Internet
System IP (system-ip)	10.254.61.61
Site ID (site-id)	120060
Hostname (host-name)	GCP-Cloud1

The following template is deployed on Chassis Number – **C8K-58FBE18E-98E3-2682-CD03-D85B43DA**

Variables	Value
Color (vpn_if_tunnel_color_value)	Biz-Internet
System IP (system-ip)	10.254.71.71
Site ID (site-id)	120060
Hostname (host-name)	GCP-Cloud2

The following template is deployed on Chassis Number – **C8K-5C361203-73F5-0E7E-3AC2-D192028EE**

© 2020 Cisco and/or its affiliates. All rights reserved.

Variables	Value
Color (vpn_if_tunnel_color_value)	Biz-Internet
System IP (system-ip)	10.254.81.81
Site ID (site-id)	120080
Hostname (host-name)	GCP-Cloud3

The following template is deployed on Chassis Number – **C8K-608A4F5C-AA71-1995-815FB1252**

Variables	Value
Color (vpn_if_tunnel_color_value)	Biz-Internet
System IP (system-ip)	10.254.91.91
Site ID (site-id)	120080
Hostname (host-name)	GCP-Cloud4

Step 13. Click **Next** to navigate to the final page.

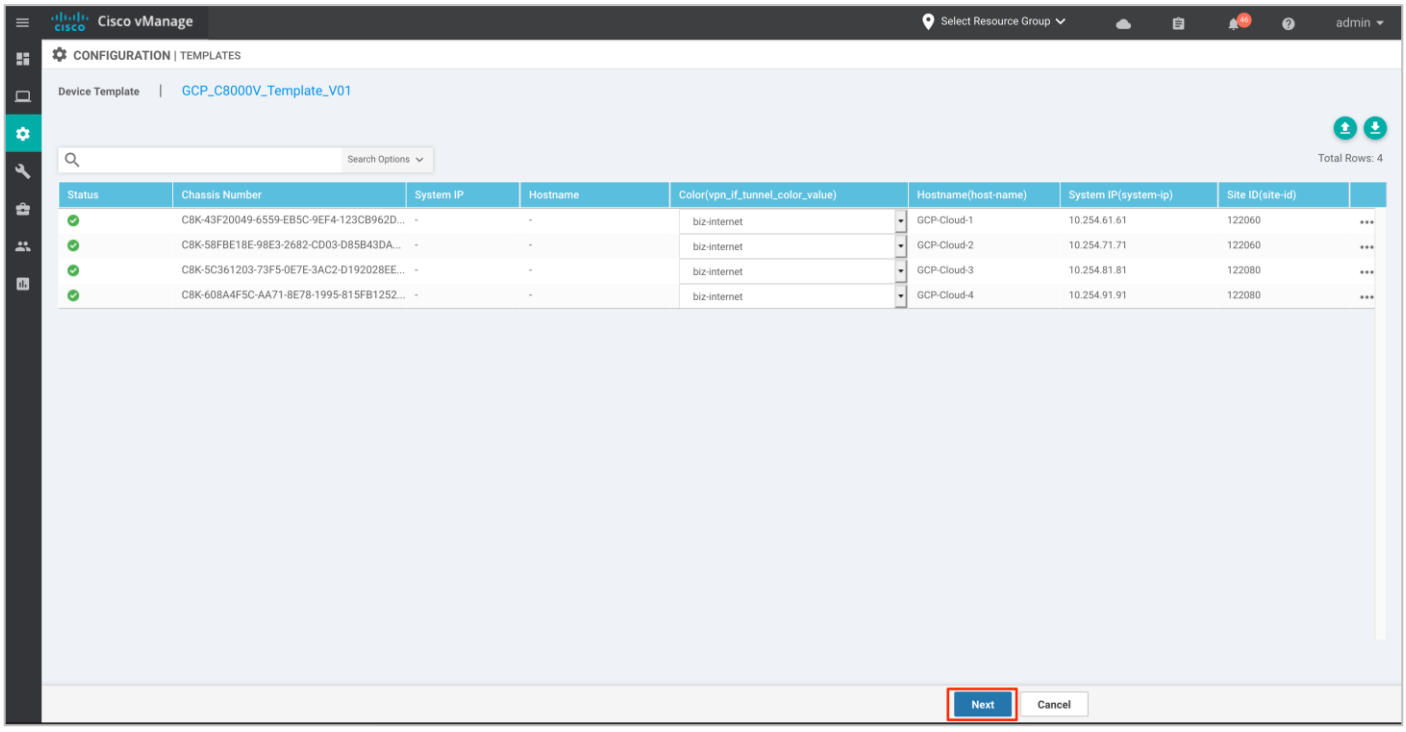

Step 14. Preview the configuration and click **Configure Devices**. Then, click **OK** to complete the template creation process.

© 2020 Cisco and/or its affiliates. All rights reserved.

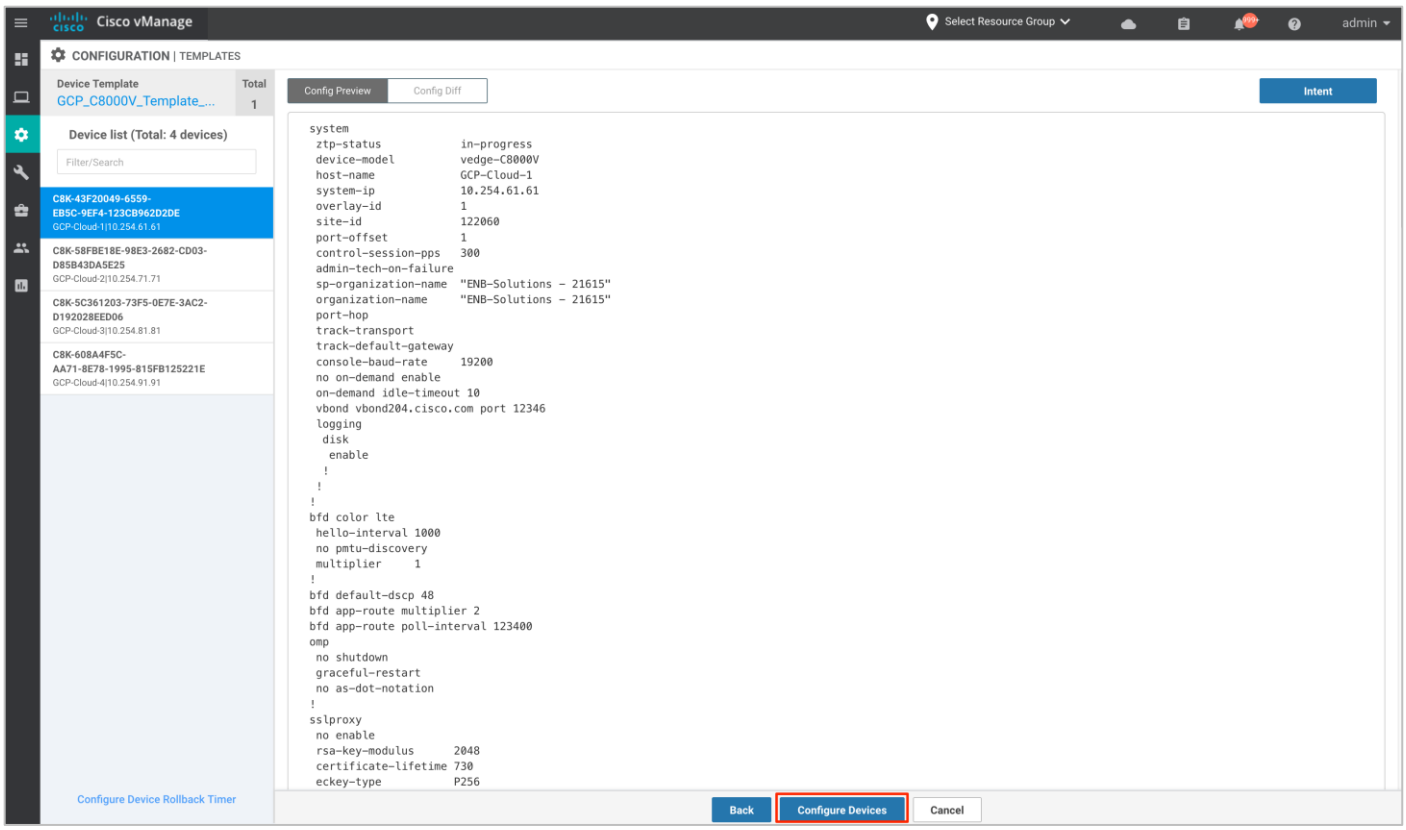

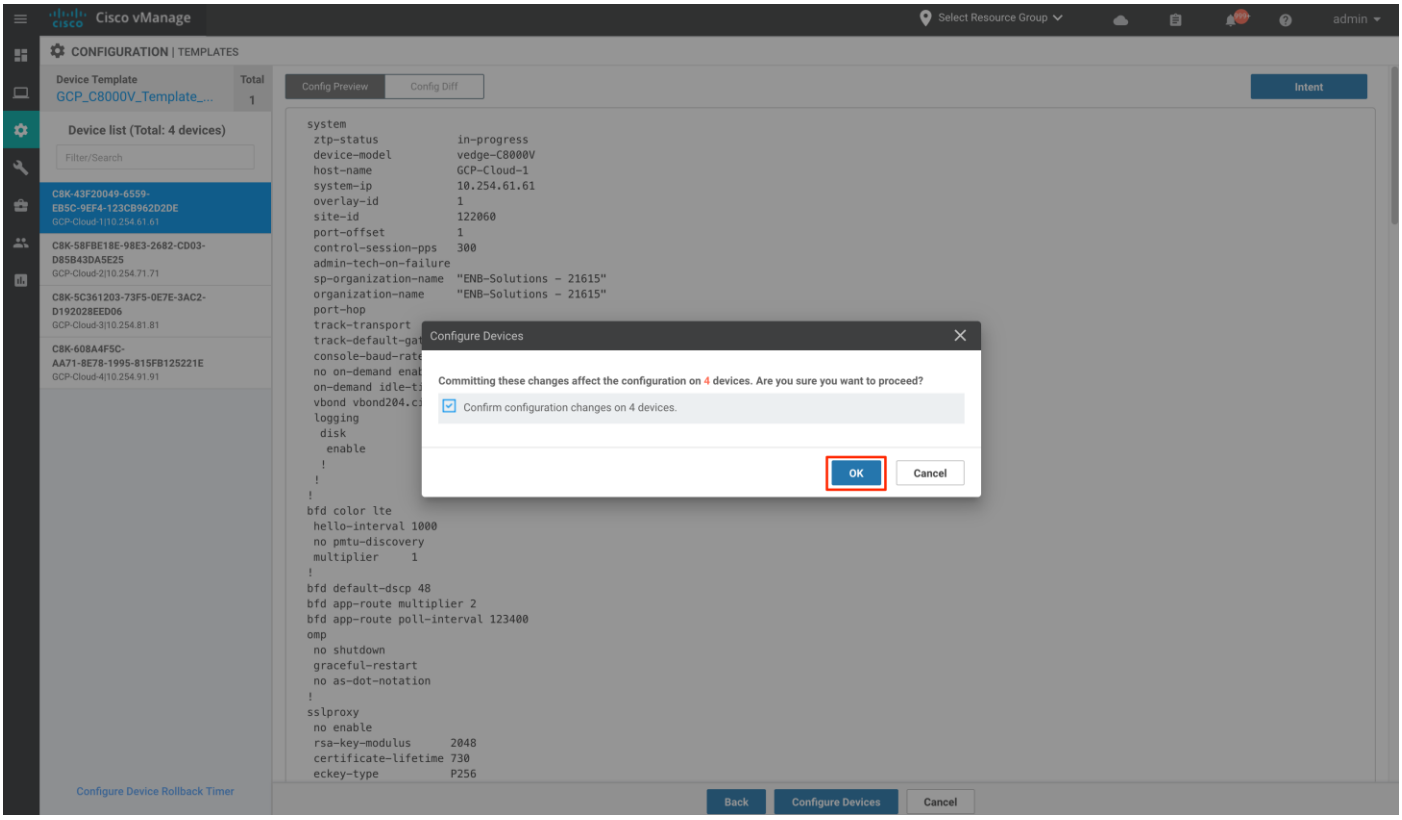

© 2020 Cisco and/or its affiliates. All rights reserved.

Process 2: Setup - Part 1 - Associate Cloud Account

This section discusses the procedures for associating your Google cloud account with Cisco vManage.

Procedure 1: Login to vManage NMS and Navigate to Cisco Cloud onRamp for Multi-Cloud

Step 1. From the navigation panel on the left side of the screen, select **Configuration > Cloud onRamp for Multi-Cloud.**

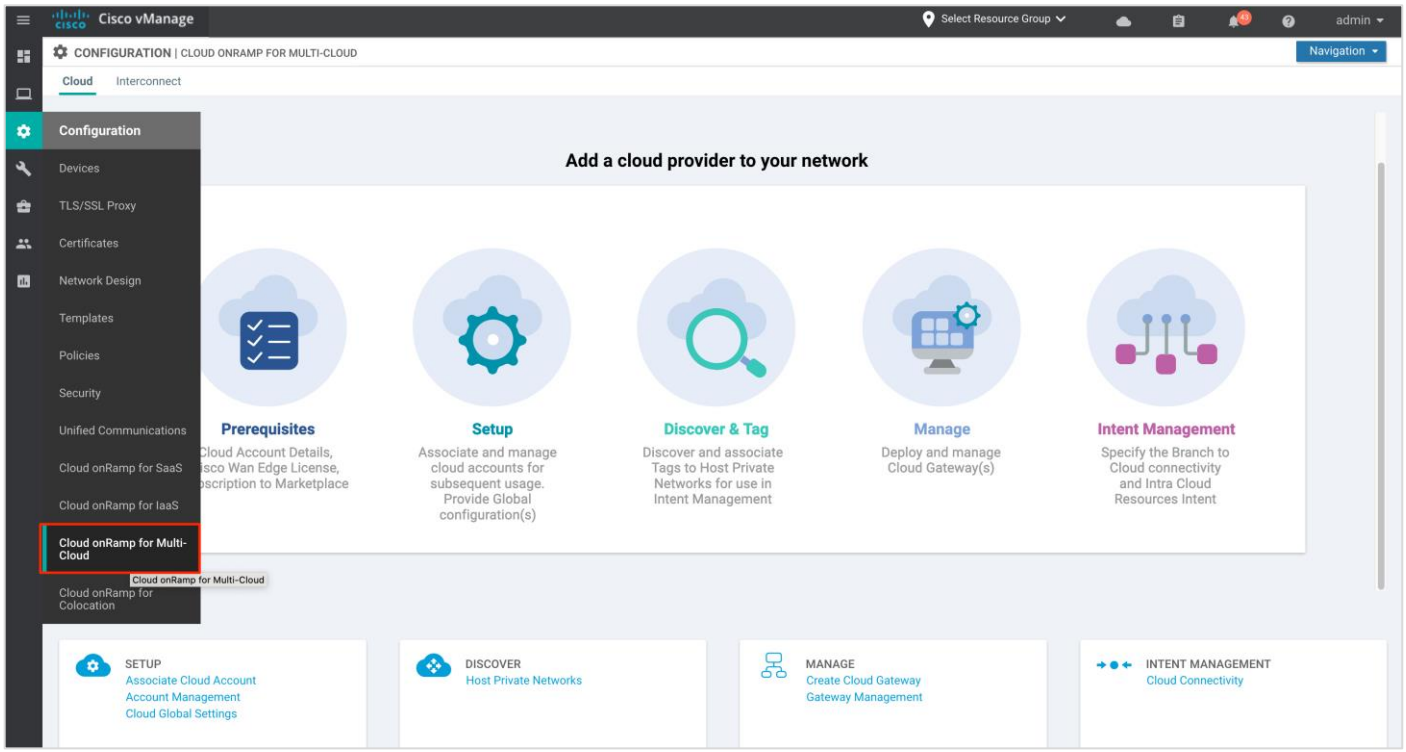

Step 2. Under **Setup**, click **Associate Cloud Account** to associate your Google cloud account with vManage.

© 2020 Cisco and/or its affiliates. All rights reserved.

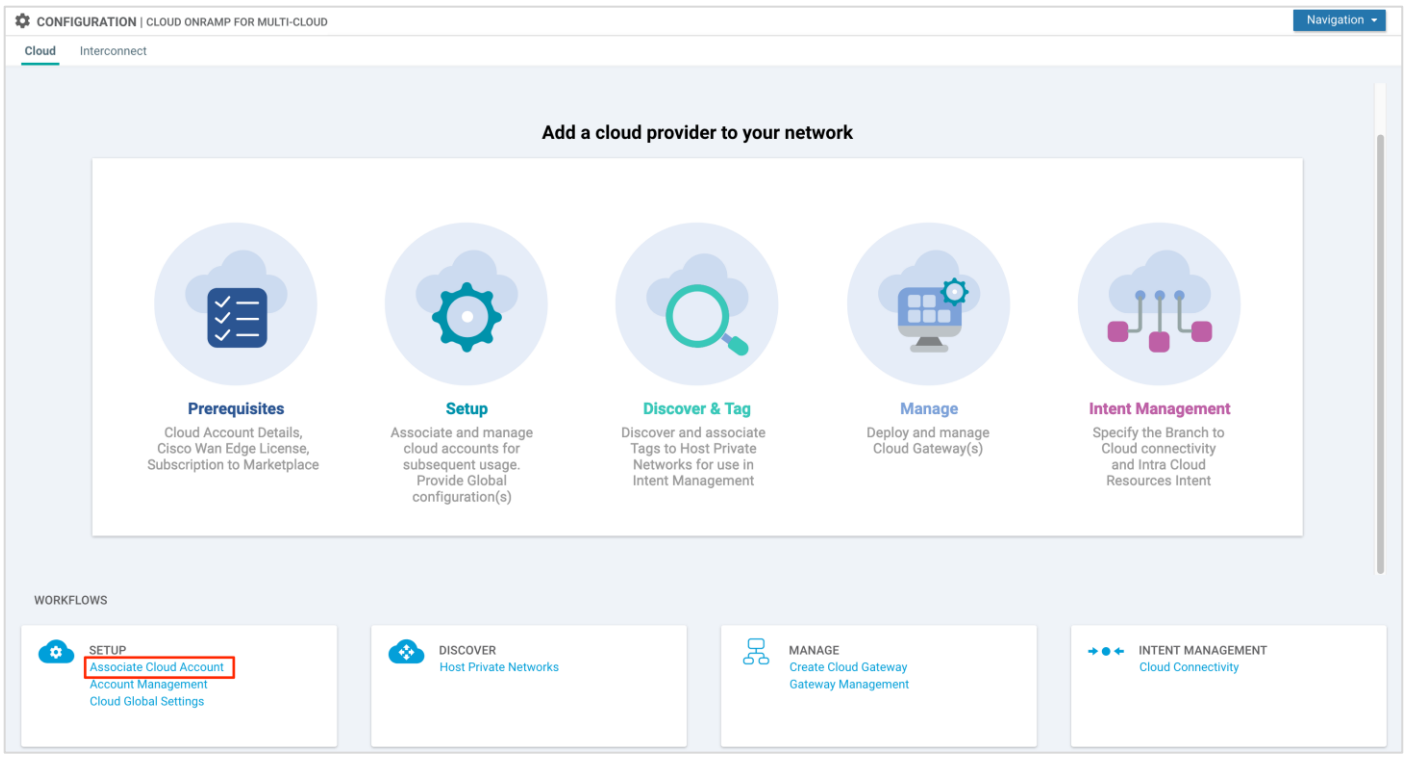

Step 3. In the **Cloud Provider** field, choose **Google Cloud** from the drop-down list and enter the requested information.

Field	Description
Cloud Account Name	Enter a name for your Google Cloud account.
Description (optional)	Enter a description for the account.
Use for Cloud Gateway	Choose **Yes** to create a cloud gateway in your account. The option **No** is chosen by default.
Billing ID (optional)	Enter the billing ID associated with your Google Cloud service account. If you provide a billing ID, it goes through an automatic billing validation process. Note: This field is visible only if you choose the option **Yes** for the **Use for Cloud Gateway** field
Private Key ID	Click **Upload Credential File** to add your Google service account's private key.

© 2020 Cisco and/or its affiliates. All rights reserved.

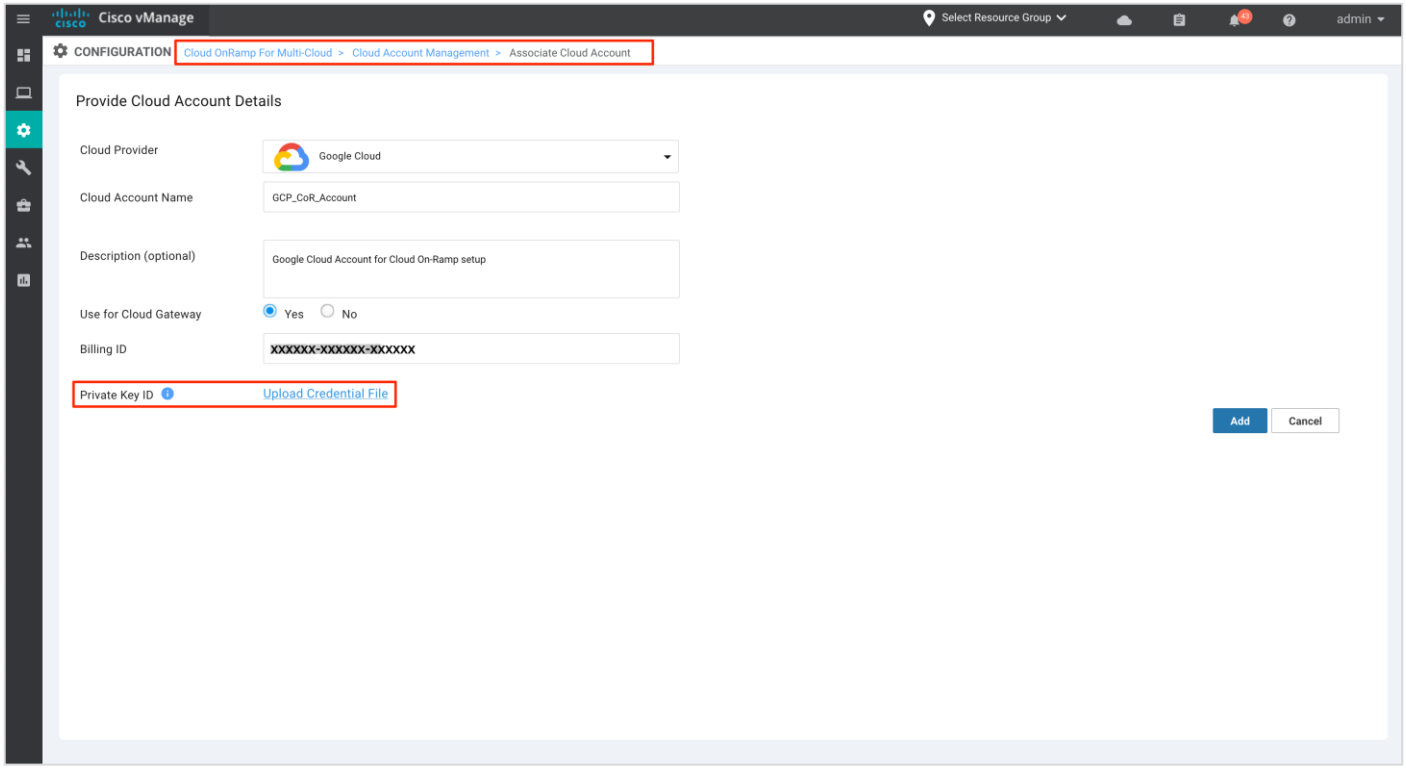

Step 4. To enter the private key, either drag and drop your Credential File or click on **Upload** to add your private key. vManage will allow user to register GCP service account keys credentials with vManage. The credentials entered will be validated in cloud and used to make GCP cloud calls.

vManage will allow two formats for entering the credentials.

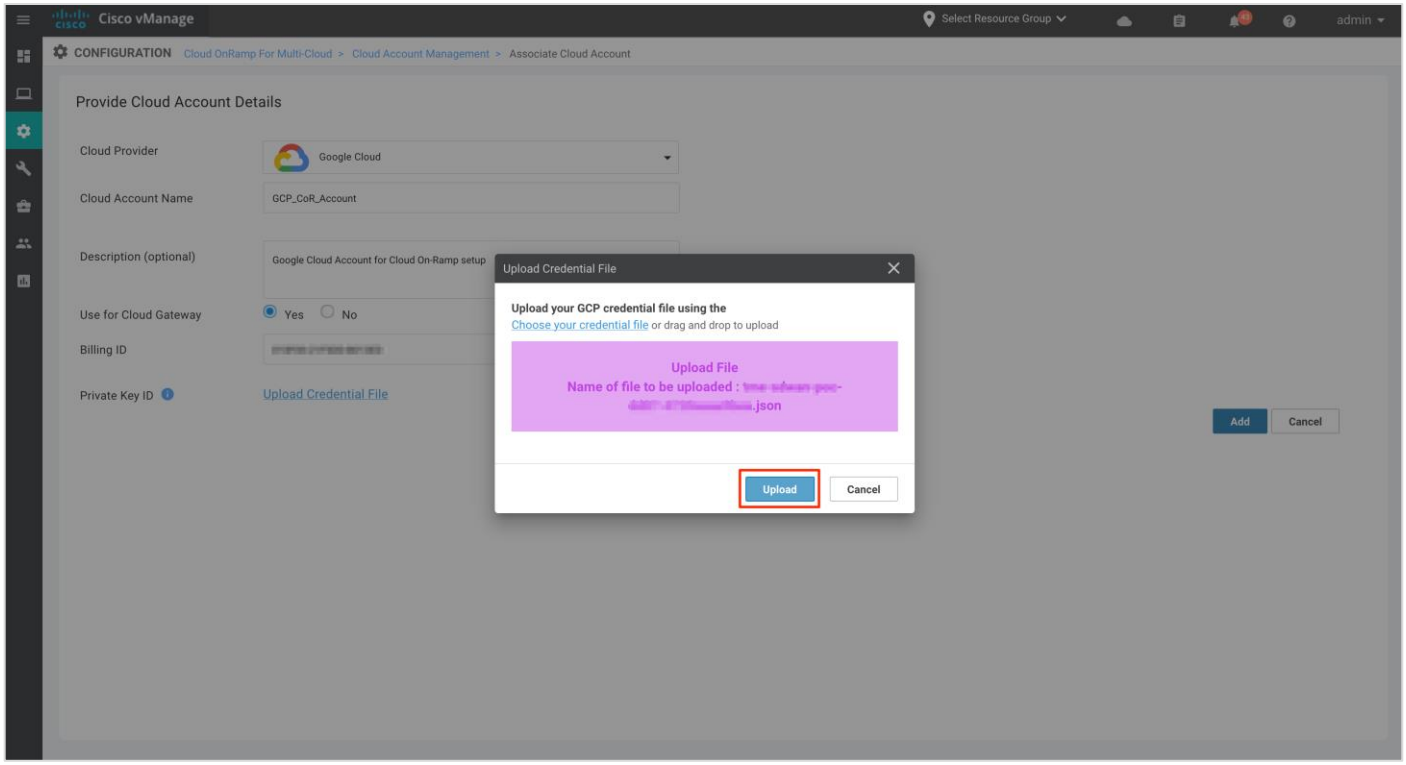

Step 5. Click **Add** to successfully associate your Google Cloud Account with vManage.

© 2020 Cisco and/or its affiliates. All rights reserved.

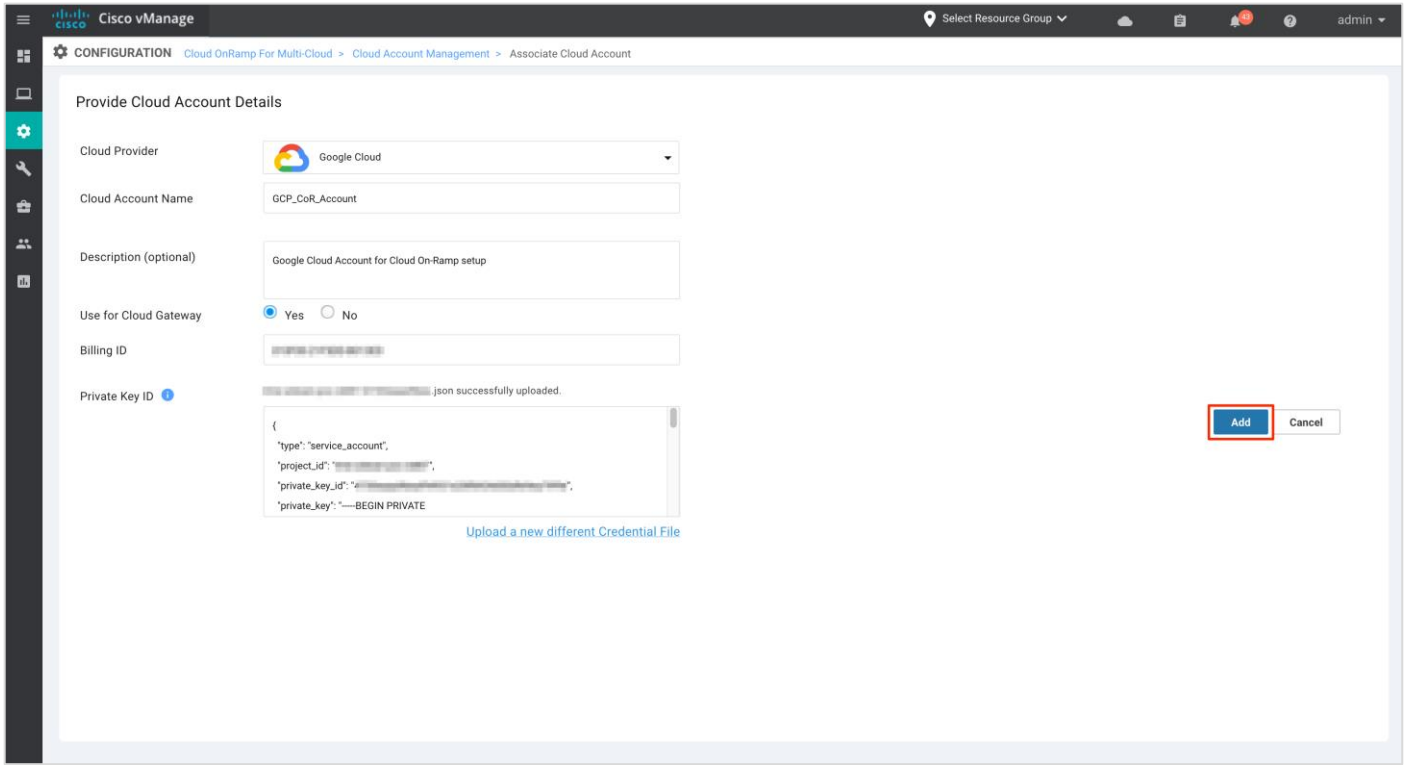

Process 2: Setup - Part 2 - Cloud Global Settings

This section discusses the steps to complete the global settings that includes the configuration items applicable system wide and used by all cloud gateway instantiations.

Step 1. Post account creation, navigate to **Cloud Global Settings**. The cloud global settings include all the global settings that apply to the cloud gateways in Google cloud.

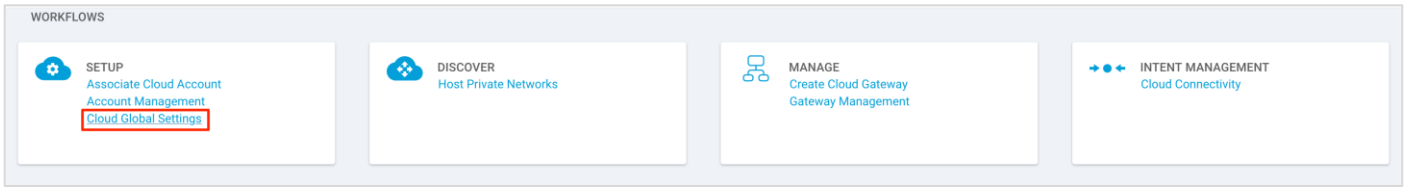

Note: You can navigate to this page from the main page that includes the workflows or click on **Navigation** located on the right corner and select **Global Settings** from the drop down.

© 2020 Cisco and/or its affiliates. All rights reserved.

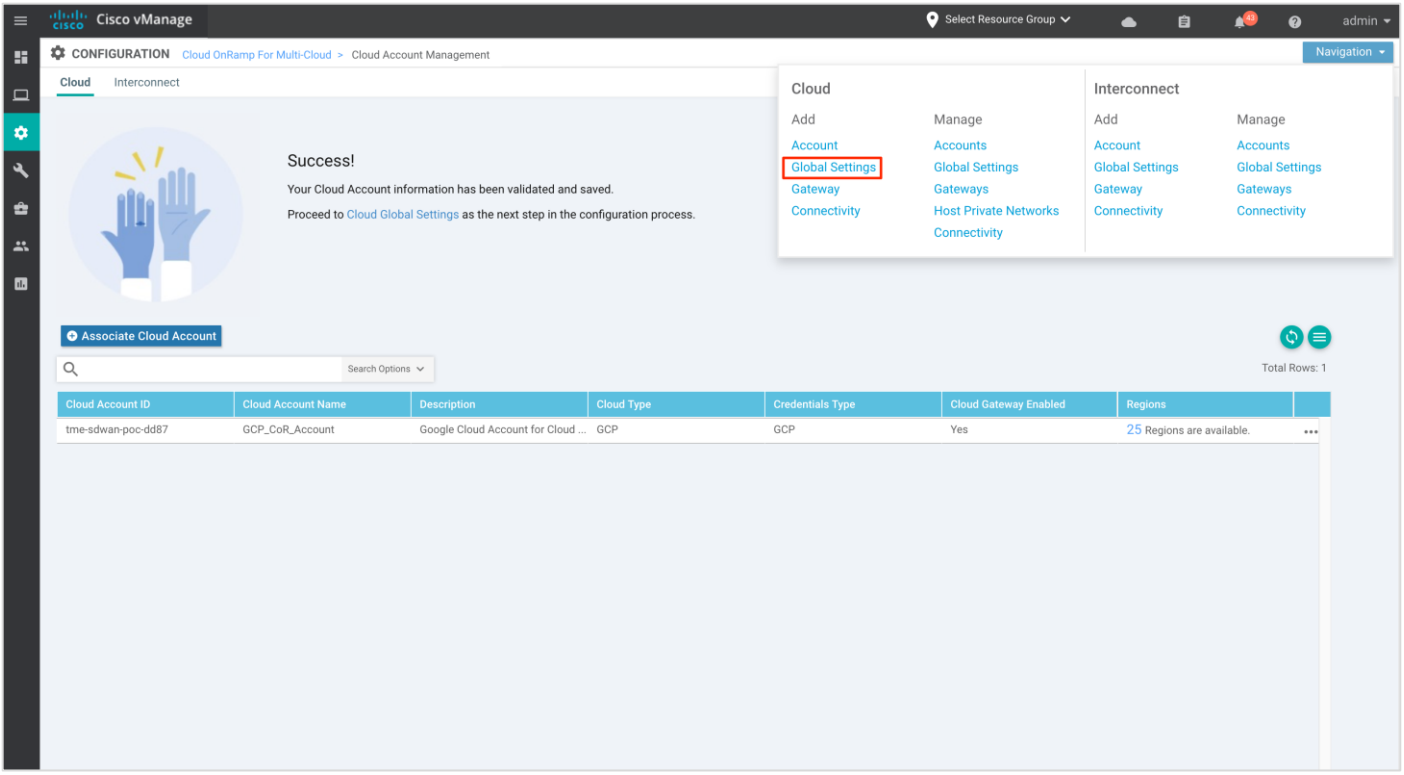

Step 2. In the **Cloud Provider** field, choose **Google Cloud** from the drop-down list and click the **Add** button to add the global cloud gateway settings. If the cloud global settings are already configured, click **Edit** to modify them.

© 2020 Cisco and/or its affiliates. All rights reserved.

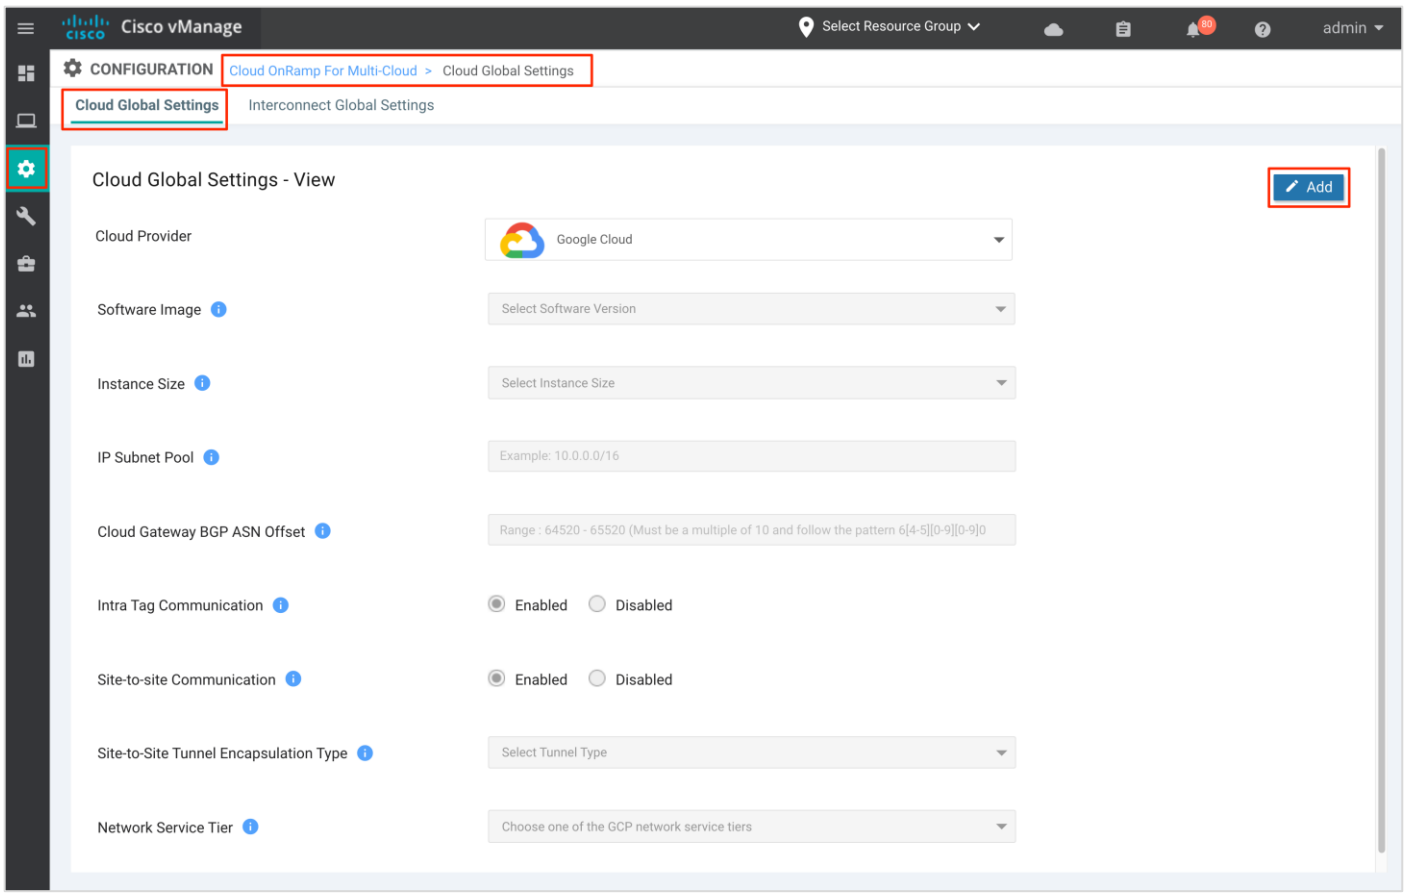

Step 3. Enter the following details to the global settings page.

Field	Details
Software Image	Choose the software image that the Cisco Catalyst80000v devices must boot up with in Google Cloud. The drop down includes Marketplace catalyst 8000v 17-04-01a and Marketplace catalyst 8000v 17-05-01a release. In this deployment, the image, Marketplace catalyst 8000v 17-05-01a is selected.
Instance Size	From the drop-down list, choose an instance based on your requirements. The drop down includes the instance sizes C4, N1-Standard-4 and C8, N1-Standard-8 instance type. Machine type N1-standard-4 provides an instance with 4 vCPUs and 15GB of memory, and N1-standard-8 provides an instance with 8 vCPUs and 30GB of memory. Note: the maximum egress bandwidth for the two machine types are different, N1-standard-4 provides 10GBPS vs 16Gbps available in N1-Standard-8 instance type.
IP Subnet Pool	Specify the IP subnet pool for all the Cisco SD-WAN cloud gateway to be deployed in Google Cloud. The subnet pool prefix must be between /16 and /21 subnet. In this deployment, the IP subnet pool is set to 10.52.0.0/16.
Cloud Gateway BGP ASN	Enter a BGP autonomous system number (ASN) between the range 64250 to 65520. • The BGP ASN entered here will be the value for the Google Cloud Routers (GCRs) in the Site-to-Cloud Transit VPC. • The entered BGP ASN plus one will be the ASN for Google Cloud Routers (GCRs) in Site-to-Site Transit VPC. • The entered ASN plus 10 will be the BGP ASN assigned to all your Catt 8000v devices.

© 2020 Cisco and/or its affiliates. All rights reserved.

	All ASNs must belong to the private ASN space.
Intra Tag Communication	Click on either of the radio button to enable or disable communication between the host VPCs part of the same tag.
Site-to-Site Communication	To enable Site-to-Site (S2S) use case select the **Enabled** radio button to establish your site-to-site transit connectivity using Google global network. Otherwise, choose Disabled. In this deployment we have set this to **Enabled** state.
Site-to-Site Tunnel Encapsulation Type	To establish the site to site communication, a second VPN 0 WAN Tunnel interface is created in the Catalyst8000v devices. The tunnel encapsulation type between the two sites can be either of type GRE or IPsec. This tab lists a drop-down that includes both GRE and IPsec encapsulation, choose one or the other. Note: The overall IPsec tunnel bandwidth between the sites cannot go beyond 3 GB, therefore Google recommends the use of GRE as it provides twice the bandwidth/throughput. Also, note regardless of the encapsulation type chosen all data traffic traversing through the Google backbone is encrypted.
Network Service Tier	Choose one of the Google Cloud service tiers. PREMIUM: Provides high-performing network experience using Google global network. STANDARD: Allows control over network costs.

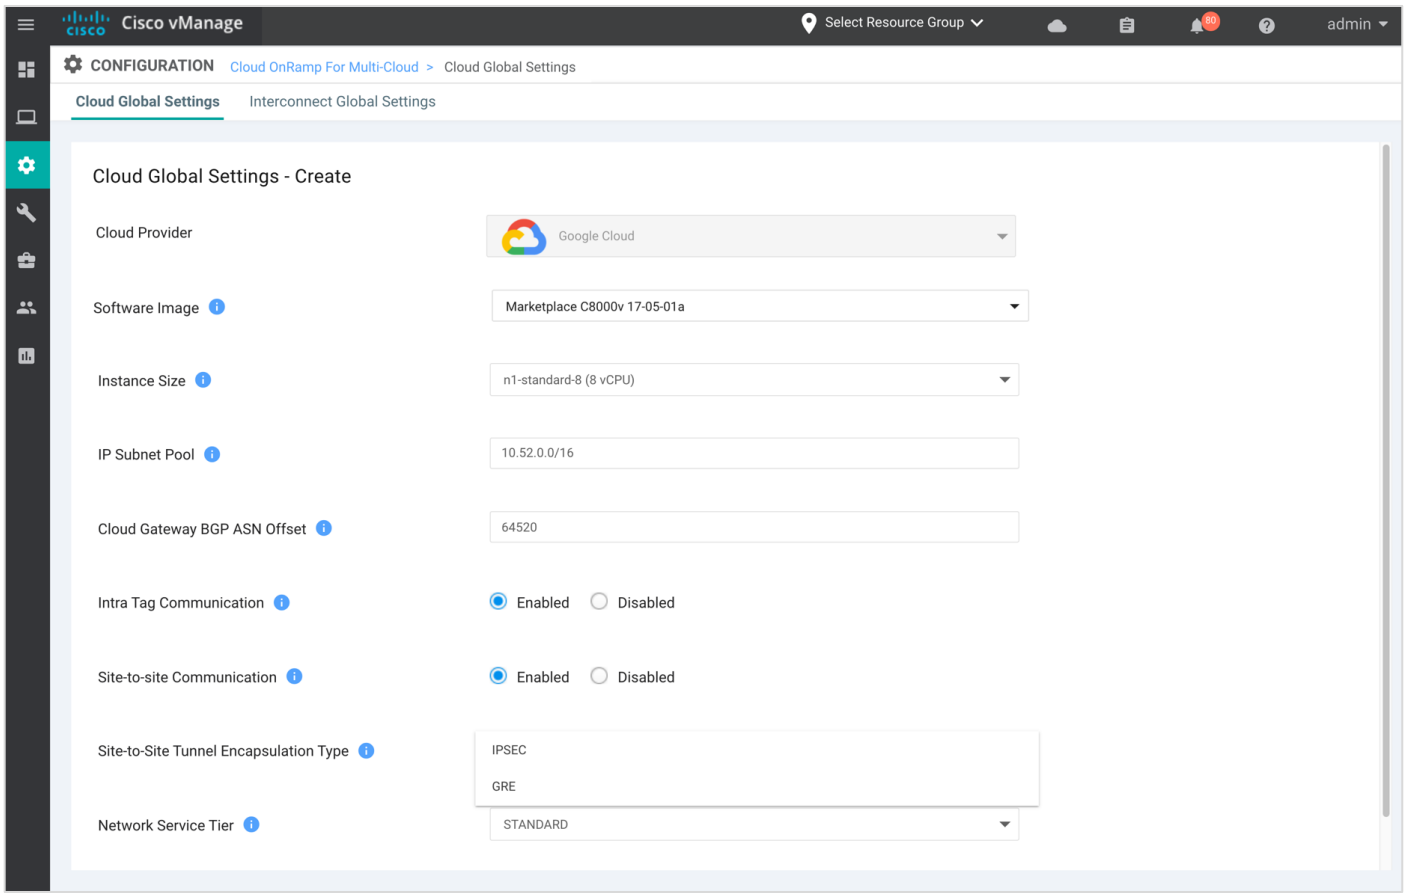

Step 4. Click **Save** or **Update** to complete the setup workflow.

© 2020 Cisco and/or its affiliates. All rights reserved.

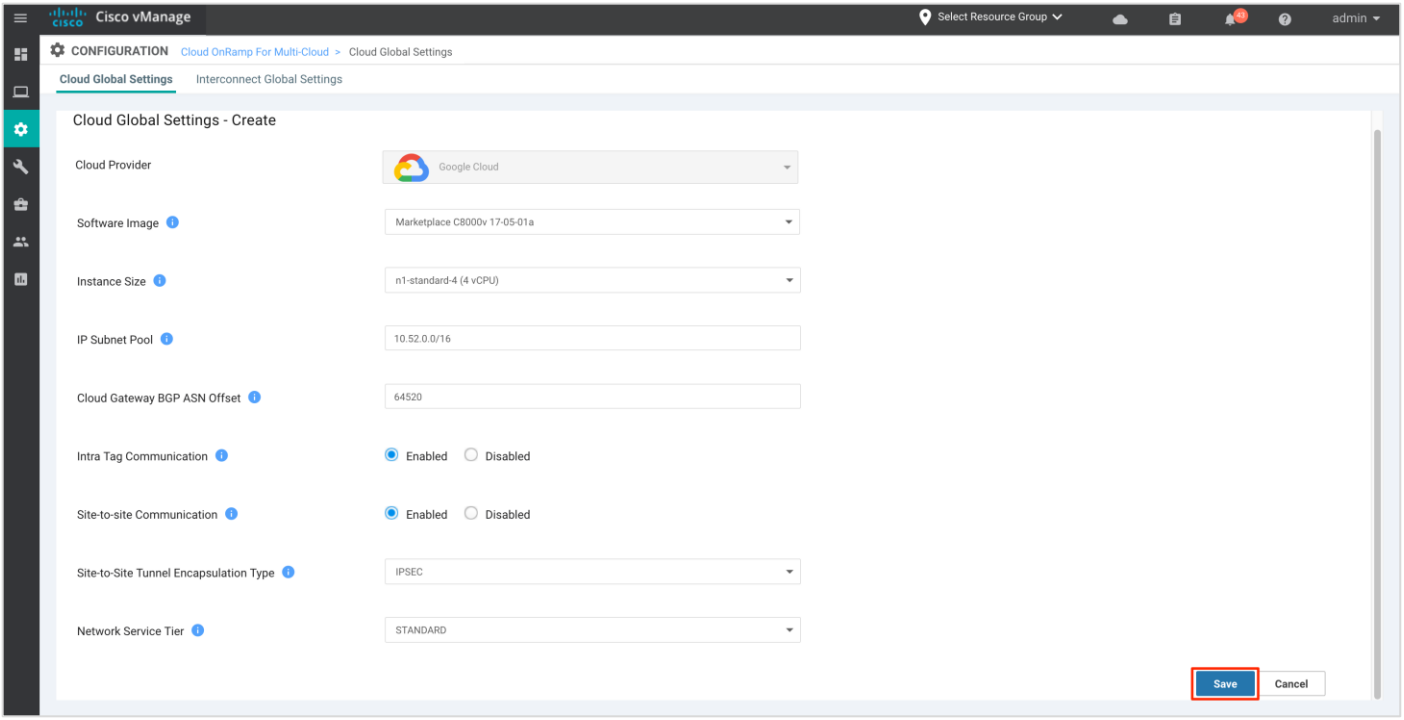

Process 3: Discover and Tags - Create Tags for the Discovered Host VPCs

After you associate your Google cloud account with Cisco vManage, you can discover your host VPCs in the regions associated with your Google cloud account. The discovered VPCs are associated to tags and these tags are used later during within Intent matrix to establish connectivity between your Google host VPCs and SD-WAN branch VPNs.

Note: If Intra Tag communication is enabled from the global settings, then all VPCs added within the same tag can communicate with each other.

Step 1. From the success page, click on **Discover Host Private Networks** to associate your host VPCs to tags.

© 2020 Cisco and/or its affiliates. All rights reserved.

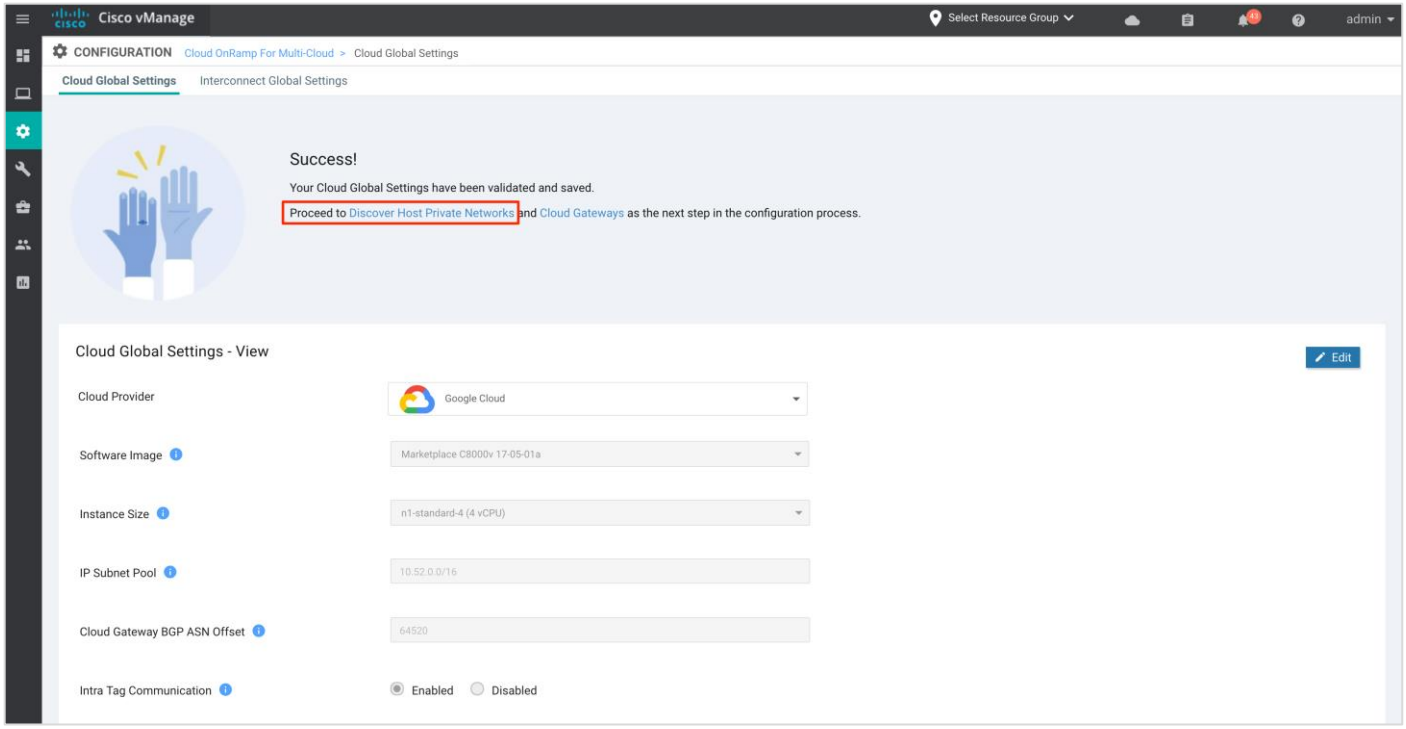

Note: Alternatively, you can also Discover host VPCs by navigating to main Cloud onRamp multi-cloud page workflow and click **Host Private Networks**.

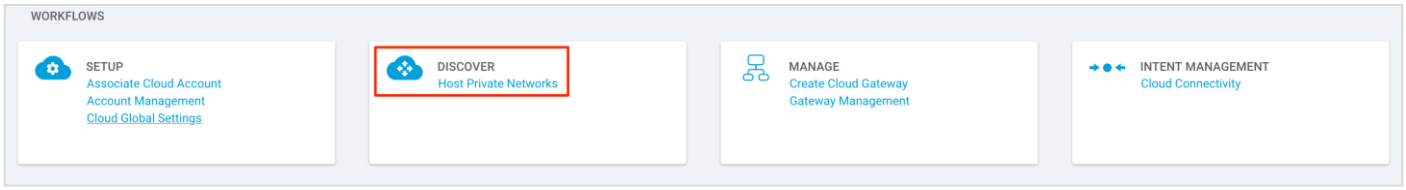

Step 2. In the **Cloud Provider** field, choose **Google Cloud**. A list of discovered Host VPCs displays in a table with the following columns: **Cloud Region**, **Account Name**, **Host VPC Name**, **Host VPC Tag**, **Account ID**, and **Host VPC ID**.

© 2020 Cisco and/or its affiliates. All rights reserved.

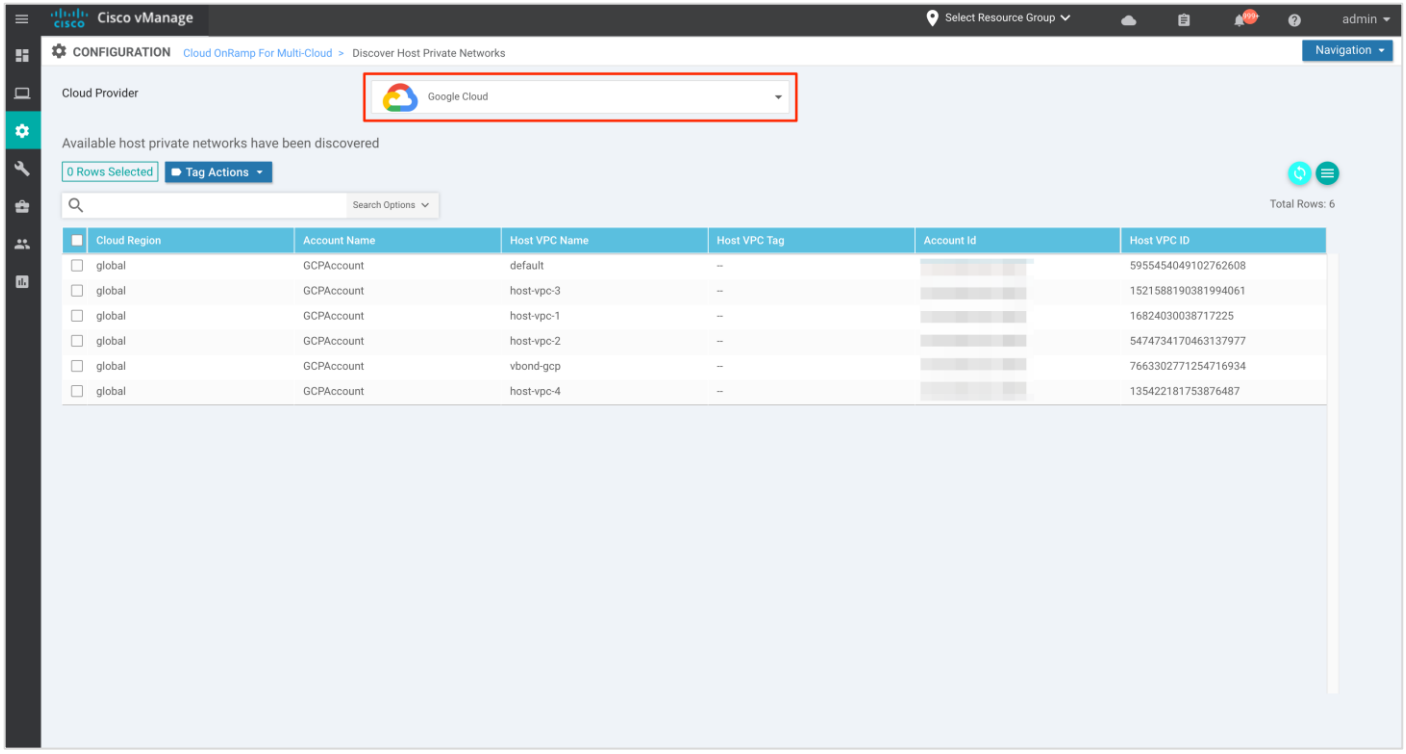

Step 3. Click the **Tag Actions** drop-down list to do any of the following:

Add Tag: Create a tag for a VPC or a group of VPCs.

Edit Tag: Change the selected VPCs for an existing tag.

Delete Tag: Delete the tag for the selected VPC.

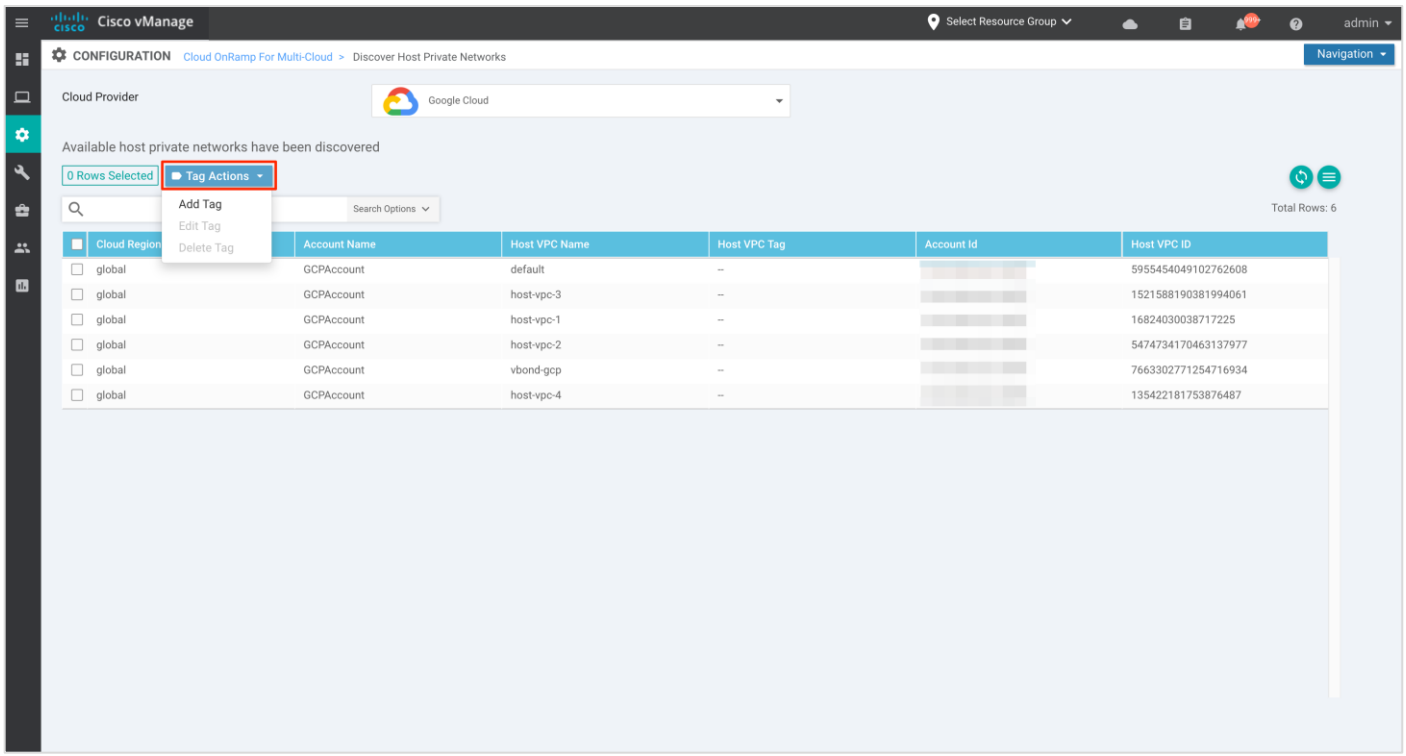

© 2020 Cisco and/or its affiliates. All rights reserved.

Step 4. Under **Add New Tag**, enter a name for the **Tag Name**, set the **Region** as **Global** and choose the host VPCs to be associated with the new tag. You can choose one or more VPCs and associate it to the same tag.

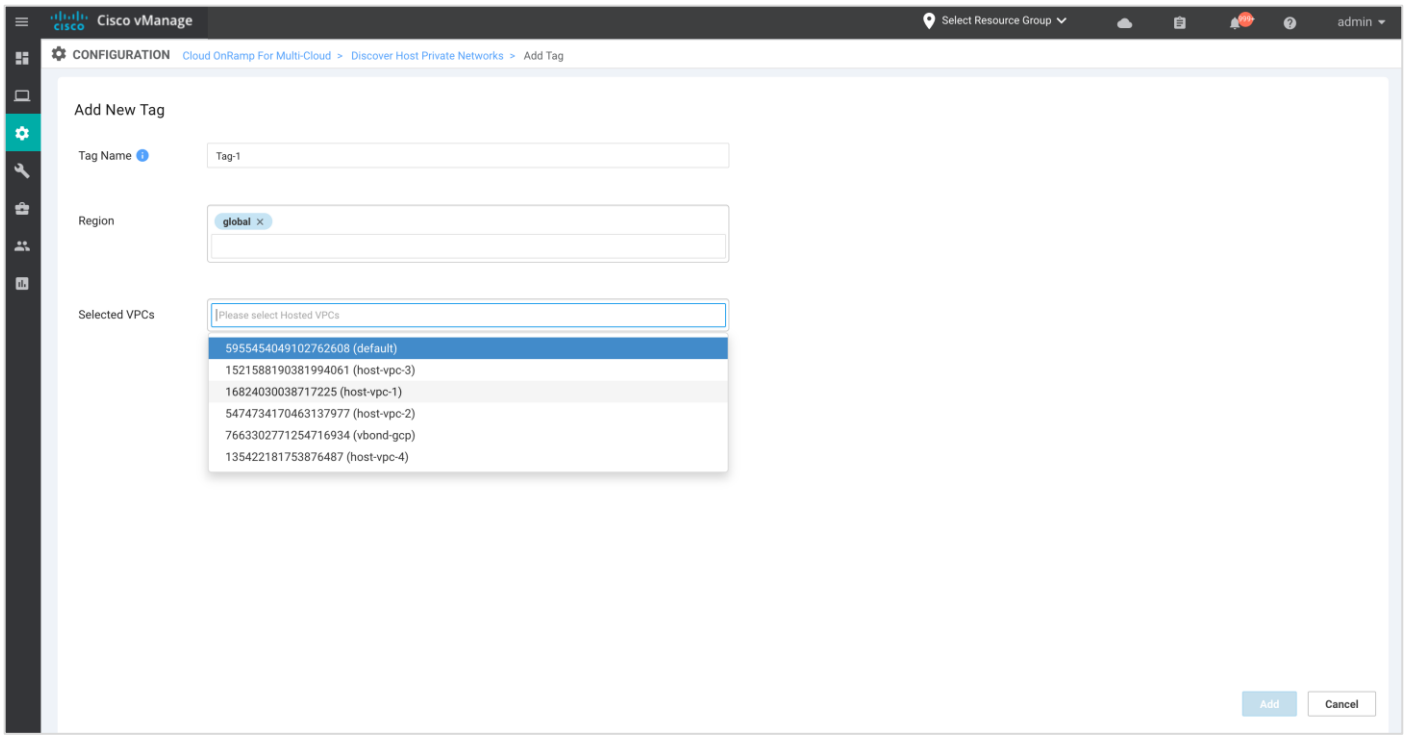

Step 5. Click **Add** to create the Tag.

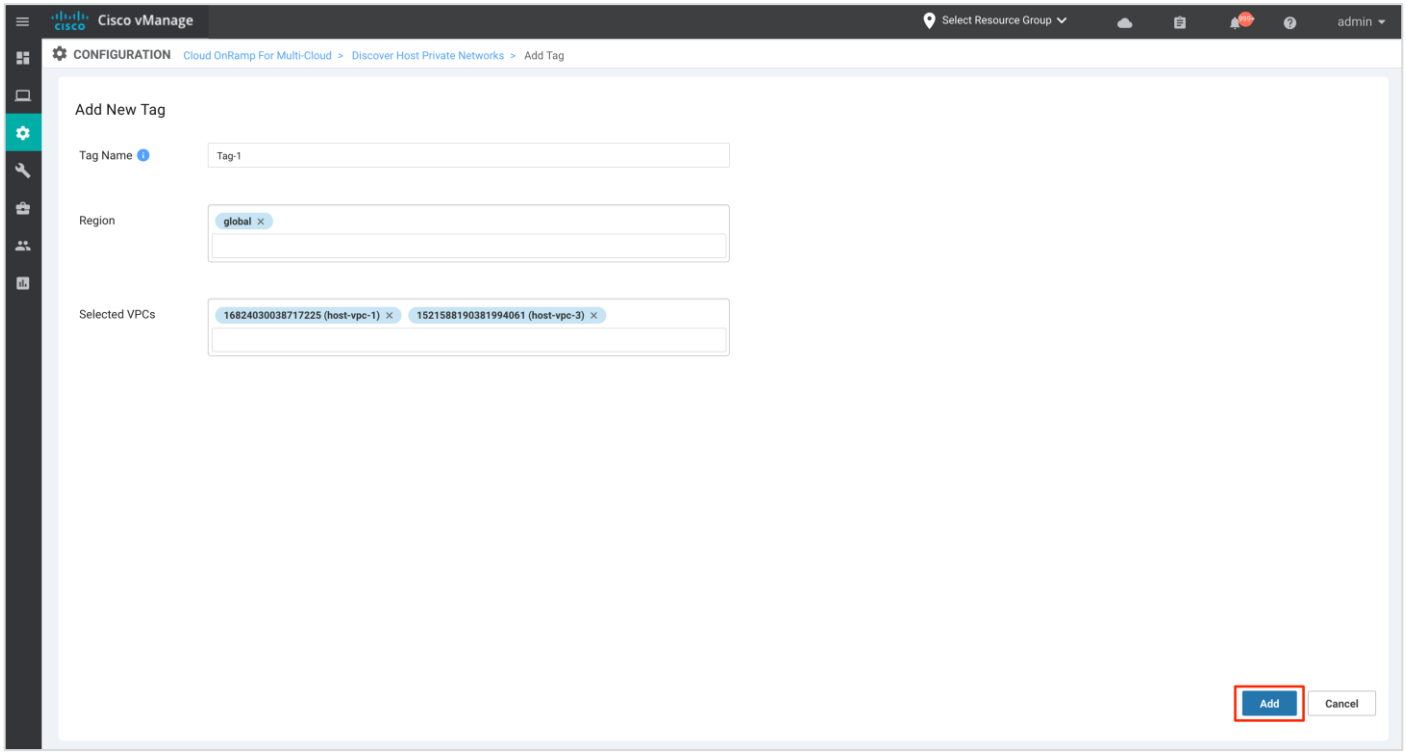

© 2020 Cisco and/or its affiliates. All rights reserved.

Process 4: Manage – Create Cloud Gateway

Once the cloud gateway global settings are added and host VPCs are associated to the tags, the next process is to instantiate a pair of Cisco Catalyst 8000V instances within the cloud gateway anchored in the three reserved VPCs –WAN transit VPC, site-to-site transit VPC, and site-to-cloud transit VPC.

Procedure 1: Create and Manage Cloud Gateways

This procedure describes how to create a Cisco SD-WAN cloud gateway with Google Cloud.

Step 1. In Cisco vManage, choose **Configuration** > **Cloud OnRamp for Multi-Cloud**.

Step 2. Under the **Manage** Workflow, click **Create Cloud Gateway**.

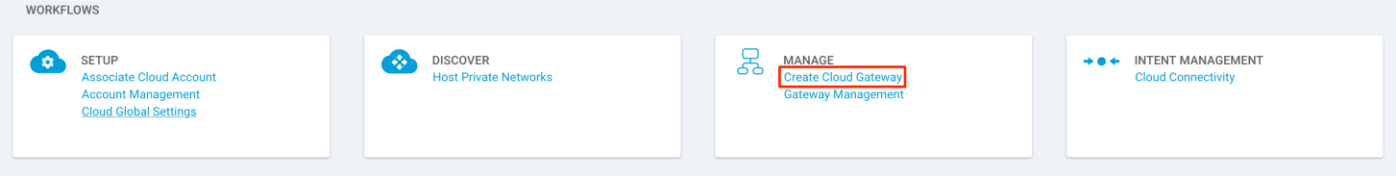

Step 3. Within the Manage Cloud Gateway Page, enter the following details.

Field		Details
Cloud Provider		Choose Google Cloud from the drop-down list
Cloud Gateway Name		Enter a name for your cloud gateway. Ensure that the name is in lowercase letters. See the Google Cloud documentation for information about naming resources and naming convention.
Description (optional)		Enter a description for the new Cloud Gateway (CGW).
Account Name		Choose your Google Cloud account name from the drop-down list.
Region		Choose a Google region from the drop-down list. The cloud gateway containing a pair of Catalyst8000v devices will be hosted in the chosen Google cloud region.
Settings	You can use either the cloud global settings or customize settings for individual Cloud Gateways (CGW) using the fields below.	
	Software Image (optional)	Edit the software image for your Catalyst8000v devices.
	Instance Size (optional)	Edit if needed the instance size for Cisco Catalyst 8000V, based on your requirements.
	IP Subnet Pool (optional)	Edit if needed the IP subnet pool to be used for the Google Cloud WAN VPC. This subnet pool needs prefixes between /16 and /21. The IP subnet pool must not overlap with the IP subnet pool specified in Cloud Global Settings.

© 2020 Cisco and/or its affiliates. All rights reserved.

	Network Service Tier (optional)	Edit the Google Cloud service tiers. PREMIUM: Provides high-performing network experience using Google global network. STANDARD: Allows control over network costs.
	UUID	Choose two Cisco Catalyst 8000V licenses from the drop-down list. These Catalyst8000v devices are deployed in the CGW.

After all the field are defined, Click **Add** to create a Cloud Gateway in Google cloud.

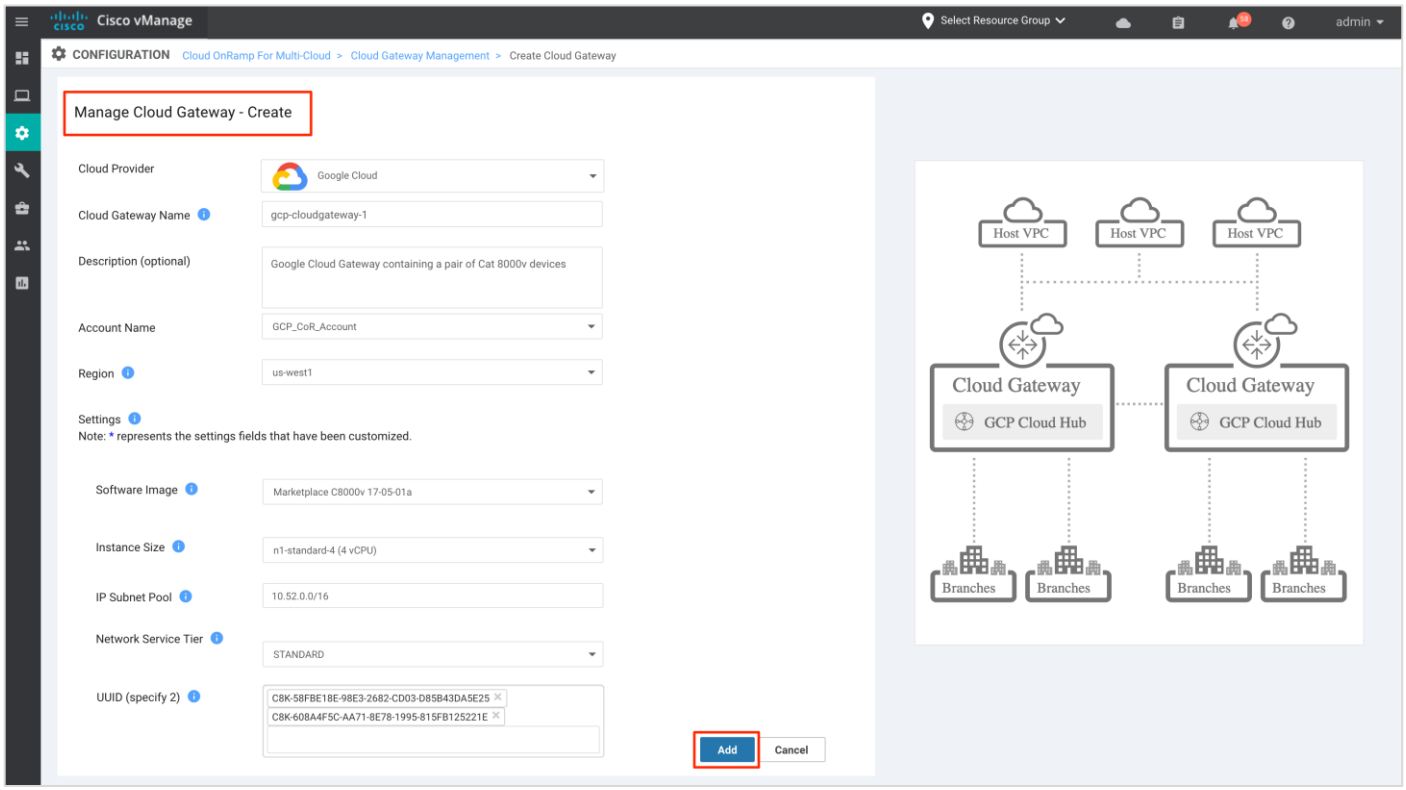

Step 4. For the Site-to-Site (S2S) use case, create a second Cloud Gateway (CGW) in a different Google cloud region. The second Cloud Gateway (CGW) also brings up another pair of Catalyst8000v in a Transit VPC that shares the common CGW resources.

Follow the steps mentioned in the step 3 to create a second cloud gateway.

Note: Only one cloud gateway can be associated in a google region.

© 2020 Cisco and/or its affiliates. All rights reserved.

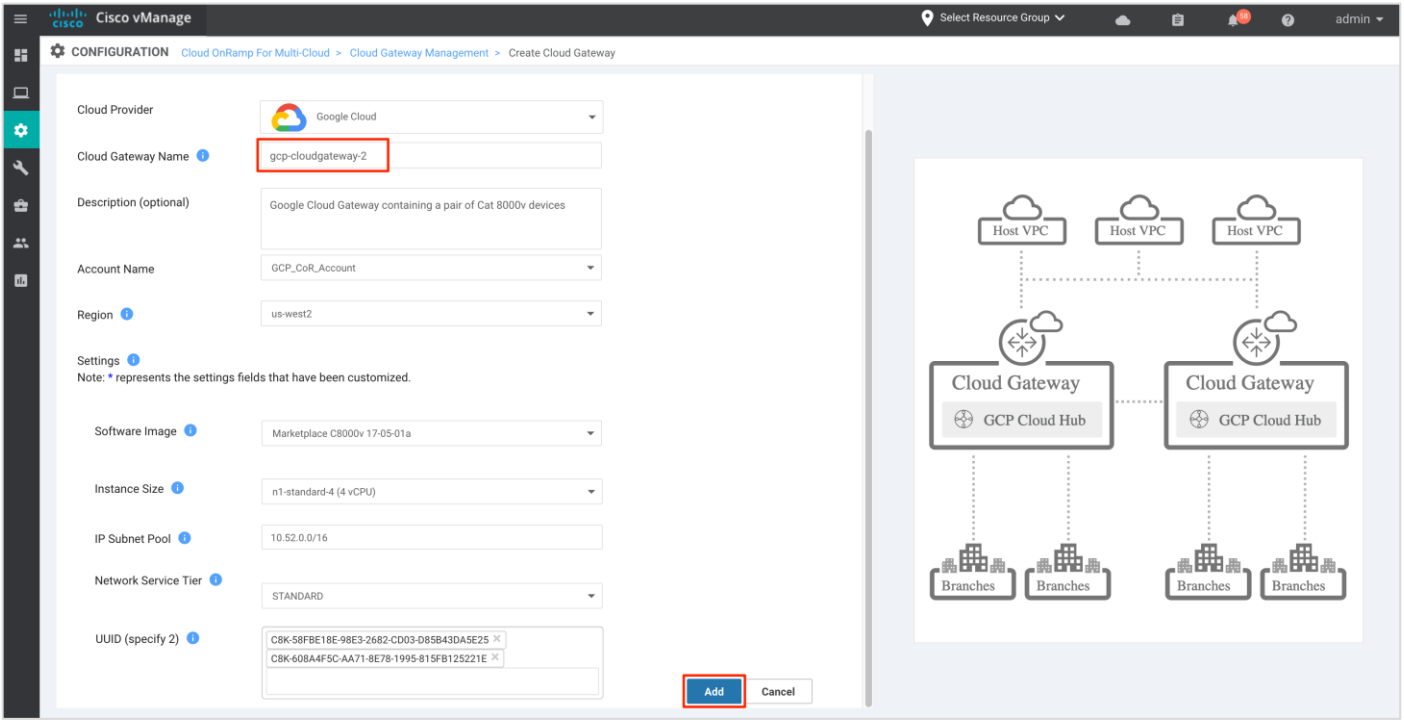

This process creates the site to cloud and site to site network connectivity (NCC) hubs, along with NCC spokes followed by site to cloud and site to site VPC, the WAN/ site to cloud/ site to site subnets are created, two cloud routers one primary and one redundant are deployed in in S2C and S2S VPC.

This ends the creation of cloud gateway; next mapping changes are identified.

Process 5: Intent Management

Only one service side VPN is created to complete the site to cloud use case. Within the intent Mapping, you can map your tags to the VPN to establish access from your site to the cloud resources.

Note: In GCP, as of the 20.5 release you can map your tags to only one VPN.

In this workflow, tag-1 and tag-2 is mapped to VPN 1.

Step 1. Navigate to **Configuration** > **Cloud OnRamp for Multi-Cloud**.

Step 2. Under **Intent Management**, click **Cloud Connectivity**. You can alternatively choose this from the navigation tab on the right.

© 2020 Cisco and/or its affiliates. All rights reserved.

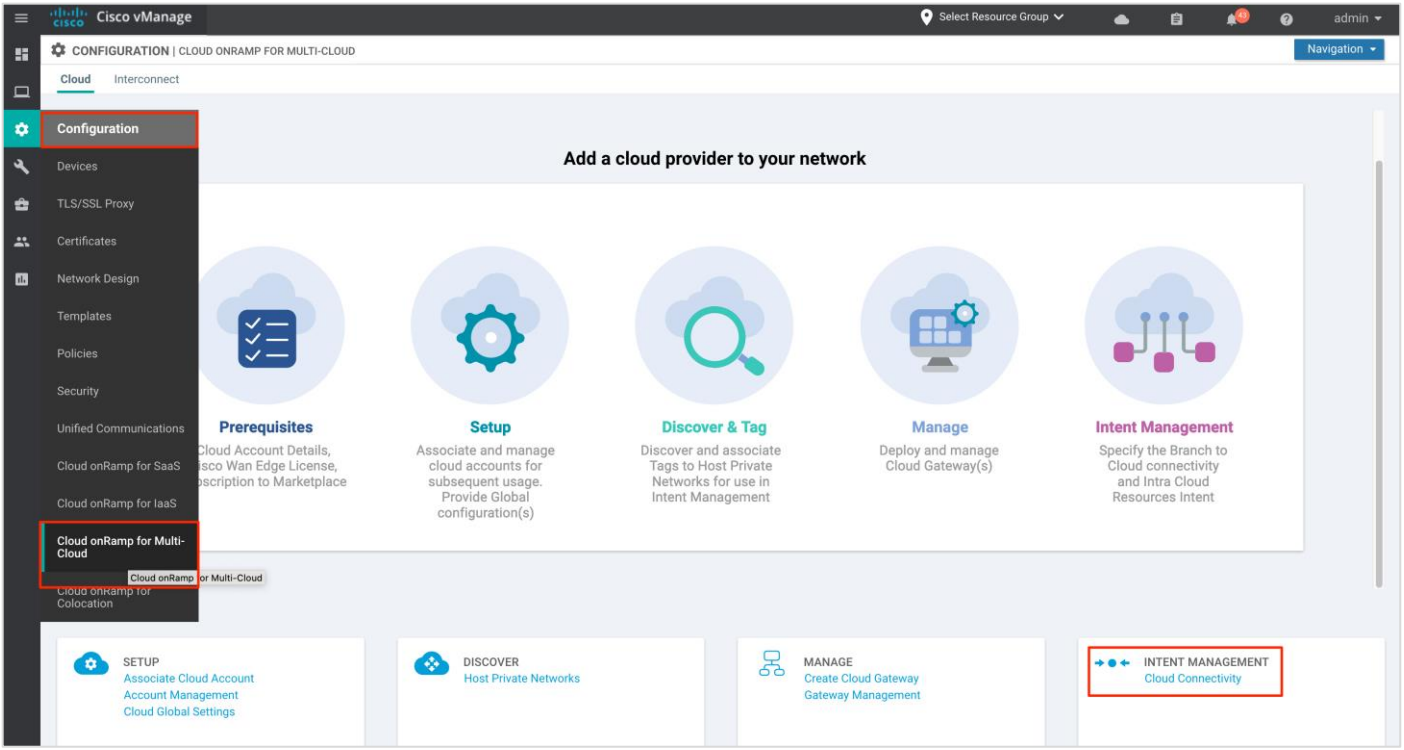

Step 3. In the **Cloud Provider** field, choose **Google Cloud** from the drop-down list.

The screen displays a connectivity matrix showing source VPNs, and their destinations. The following legend provides information about the status of the intent:

Blue: Intent Defined

Green: Intent Realized

Red: Intent Realized with Errors

Click any of the cells in the matrix to get a more detailed status information.

Under **Mapping**, define or record a new intent. Click **Edit** to map a VPN to a configured Tag.

© 2020 Cisco and/or its affiliates. All rights reserved.

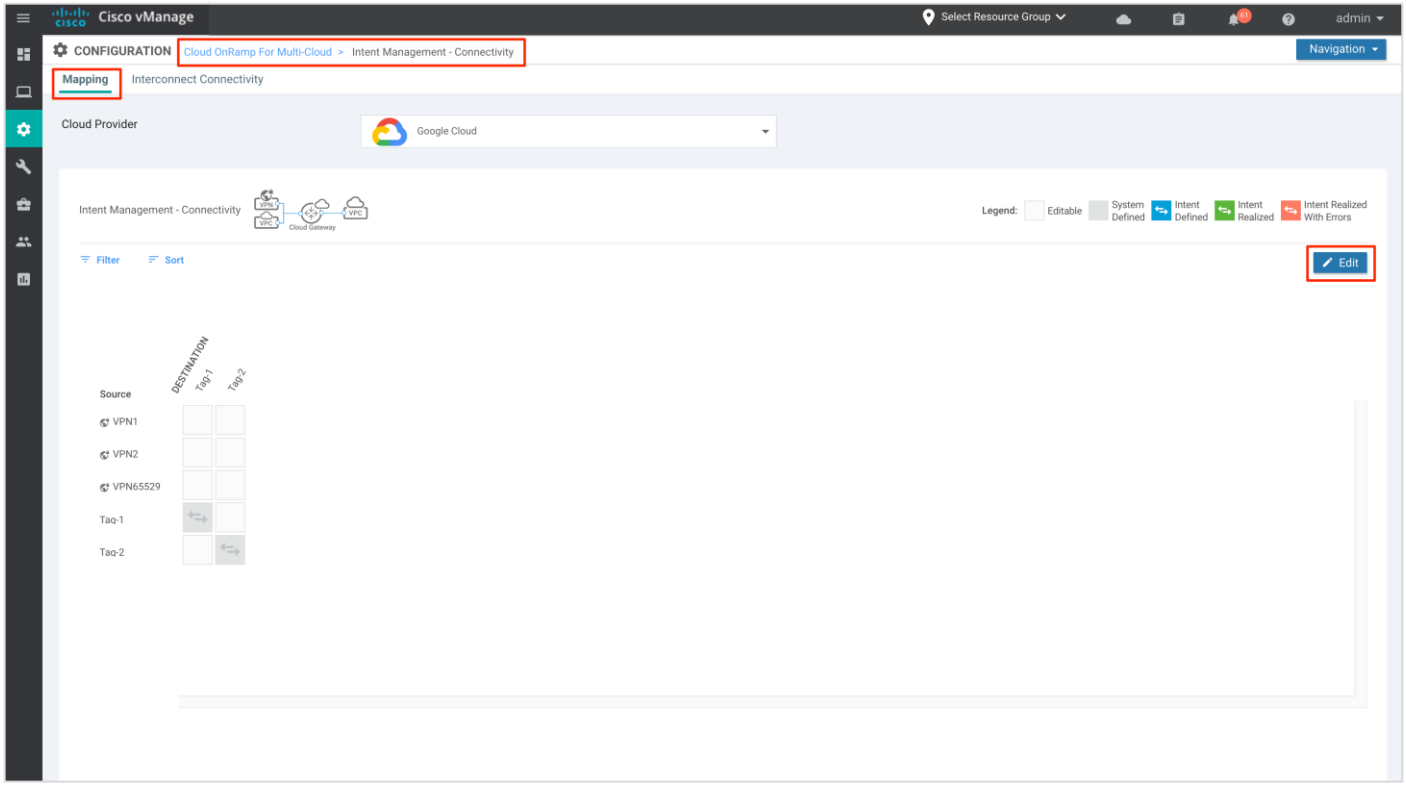

Technical Tip

Within the Intent Management – Connectivity, you might notice one or more VPN as source. The VPN list here is based on the VPN Segments available under Administration tab. Regardless of how many VPNs you view here, through the Google cloud onRamp workflow only one service-side VPN is deployed in each Catalyst 8000v for the site to cloud use case. Therefore, you can map tags to only one VPN.

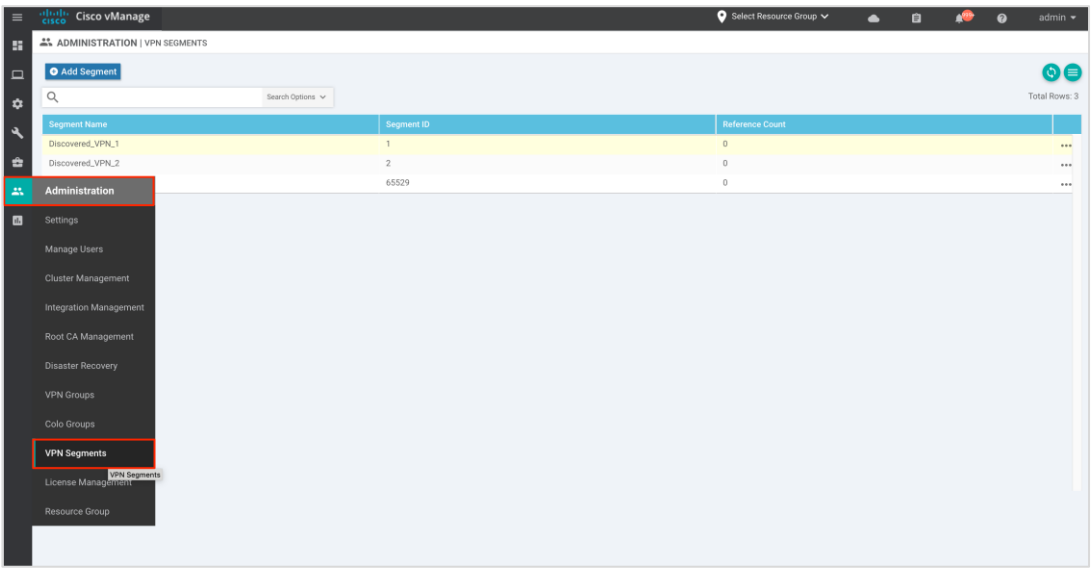

© 2020 Cisco and/or its affiliates. All rights reserved.

Step 4. Choose the cells that correspond to a VPN and the VPC tags associated with it and click **Save**. In this guide, we are mapping VPN 1 to tag 1 and tag 2 VPC workloads.

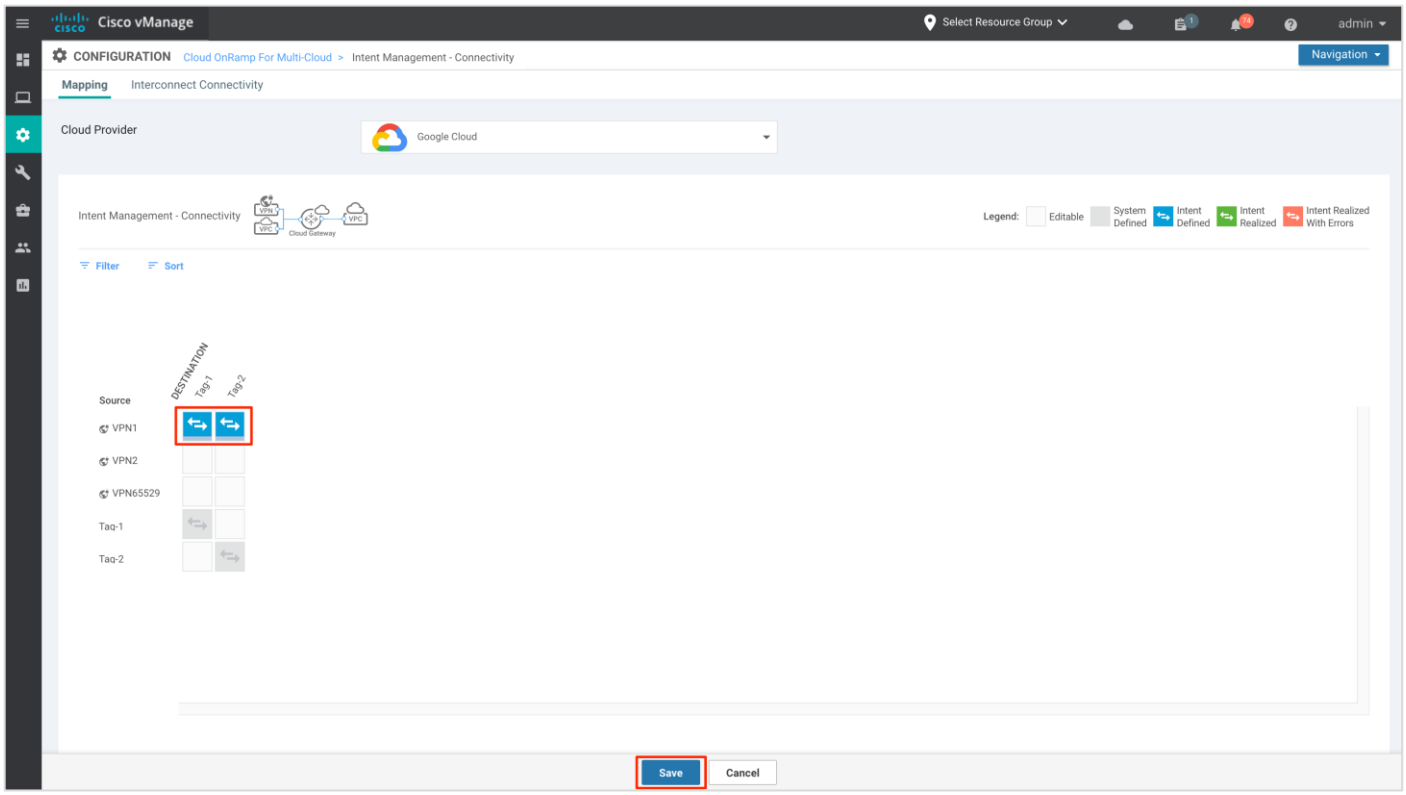

The following screenshot displays the **TASK VIEW** page upon completion of **Multicloud-Connectivity Mapping**.

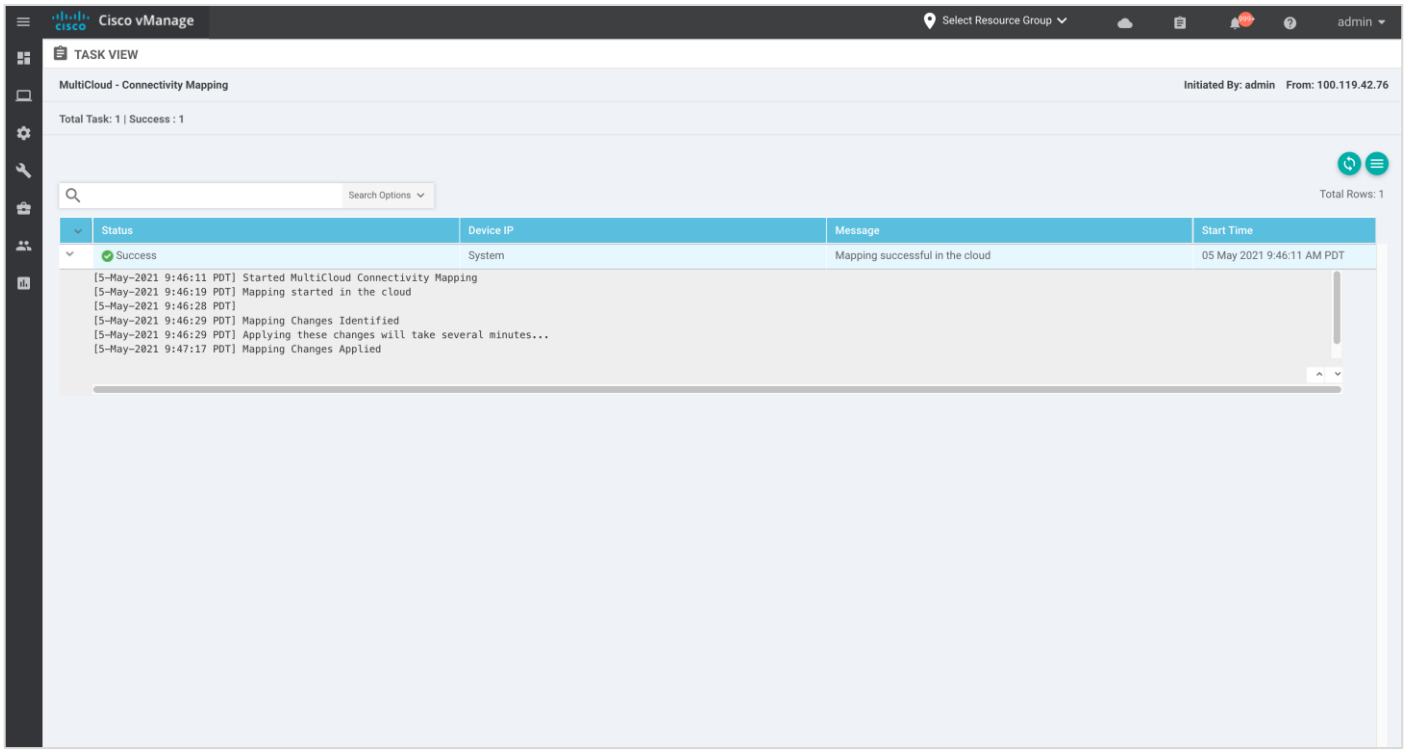

© 2020 Cisco and/or its affiliates. All rights reserved.

This completes the Cloud on-Ramp for Multicloud workflow.

Process 6: Deletion of Resources

If you want to delete and clean up resources, which were created by the Cloud onRamp GCP automation, follow the order below.

- Delete NCC Spokes using Google cloud CLI
- Delete NCC Hubs using Google cloud CLI
- Delete GCRs using GCP Console (Hybrid Connectivity -> Cloud Routers page) or CLI
- Delete c8kv VMs using GCP Console or CLI
- Delete VPCs using GCP Console or CLI

Example for Google cloud CLI listing two NCC hubs – the first hub is used for site-to-site and the second hub is for site-to-cloud use case:

```
cloudshell: ~ (gcp)$ gcloud alpha network-connectivity hubs list
[
  {
    "createTime": "2021-03-17T06:01:36.984779764Z",
    "name": "projects/gcp-XXXXX635/locations/global/hubs/s2s-ncc-hub--emo-xxxxx",
    "spokes": [

"https://networkconnectivity.googleapis.com/v1alpha1/projects/53XXXXX3/locations/austra
lia-southeast1/spokes/gc-sydney-cgw1--emo-xxxxx-australia-southeast1-s2s-ncc-spoke",

"https://networkconnectivity.googleapis.com/v1alpha1/projects/5325XXXXX/locations/us-
west2/spokes/gc-uswest-cgw1--emo-xxxxx-us-west2-s2s-ncc-spoke"
    ],
    "state": "ACTIVE",
    "uniqueId": "603ab47a-8979-XXXXX-5XXXX6b9d360e",
    "updateTime": "2021-03-17T06:01:37.281770511Z"
  },
  {
    "createTime": "2021-03-17T05:58:25.072124681Z",
    "name": "projects/gcp-npitaev20XXXX/locations/global/hubs/s2c-ncc-hub--emo-xxxxx",
    "spokes": [

"https://networkconnectivity.googleapis.com/v1alpha1/projects/5325XXXX523/locations/aus
tralia-southeast1/spokes/gc-sydney-cgw1--emo-xxxxx-australia-southeast1-s2c-ncc-spoke",

"https://networkconnectivity.googleapis.com/v1alpha1/projects/5325XXXX6523/locations/us
-west2/spokes/gc-uswest-cgw1--emo-xxxxx-us-west2-s2c-ncc-spoke"
    ],
    "state": "ACTIVE",
    "uniqueId": "4236ce9b-210f-XXXXde-7b566e39f274",
    "updateTime": "2021-03-17T05:58:25.357399310Z"
  }
```

© 2020 Cisco and/or its affiliates. All rights reserved.

]

Operate - Cisco Cloud onRamp for Multi-Cloud Monitoring

Using the vManage GUI, you can monitor, troubleshoot, and manage the Cisco SD-WAN Cloud onRamp for Multi-Cloud using Google Cloud. The 3 main ways to monitor the deployment is as given below:

vManage Cloud onRamp for Multi-Cloud Dashboard: From the vManage Cloud onRamp for Multi-Cloud dashboard you can monitor the connectivity state of the Cloud Gateway (CGW), the mapped tags and host VPCs within those tags.

vManage Monitor Dashboard: From the vManage Monitor dashboard, you can view and gather error logs and interface details for the Catalyst 8000v routers deployed in the CGW.

vManage SSH Dashboard: From the vManage SSH dashboard, you can view the CLI configs and monitor route updates using show commands for all the Catalyst 8000v devices deployed in Google Clouds CGWs.

Process 1: vManage Cisco Cloud onRamp for Multi-Cloud Dashboard

Procedure 1: Reachability to Cisco Catalyst 8000v Instances

Verify the bring-up and reachability of the Cisco Catalyst 8000v instances provisioned within the cloud gateway(s).

Step 1. In Cisco vManage, choose **Configuration** > **Cloud onRamp for Multi-Cloud**.

Step 2. Under the **Cloud** section located at the top of the screen, the **Network Snapshot** displays a summary of the cloud gateways deployed, the number of Google host VPCs connected to the Site via Tags, along with its Site-Cloud IP address and VPN label mapped to, finally followed by the status of Catalyst 8000v edge devices brought up to form the cloud gateway.

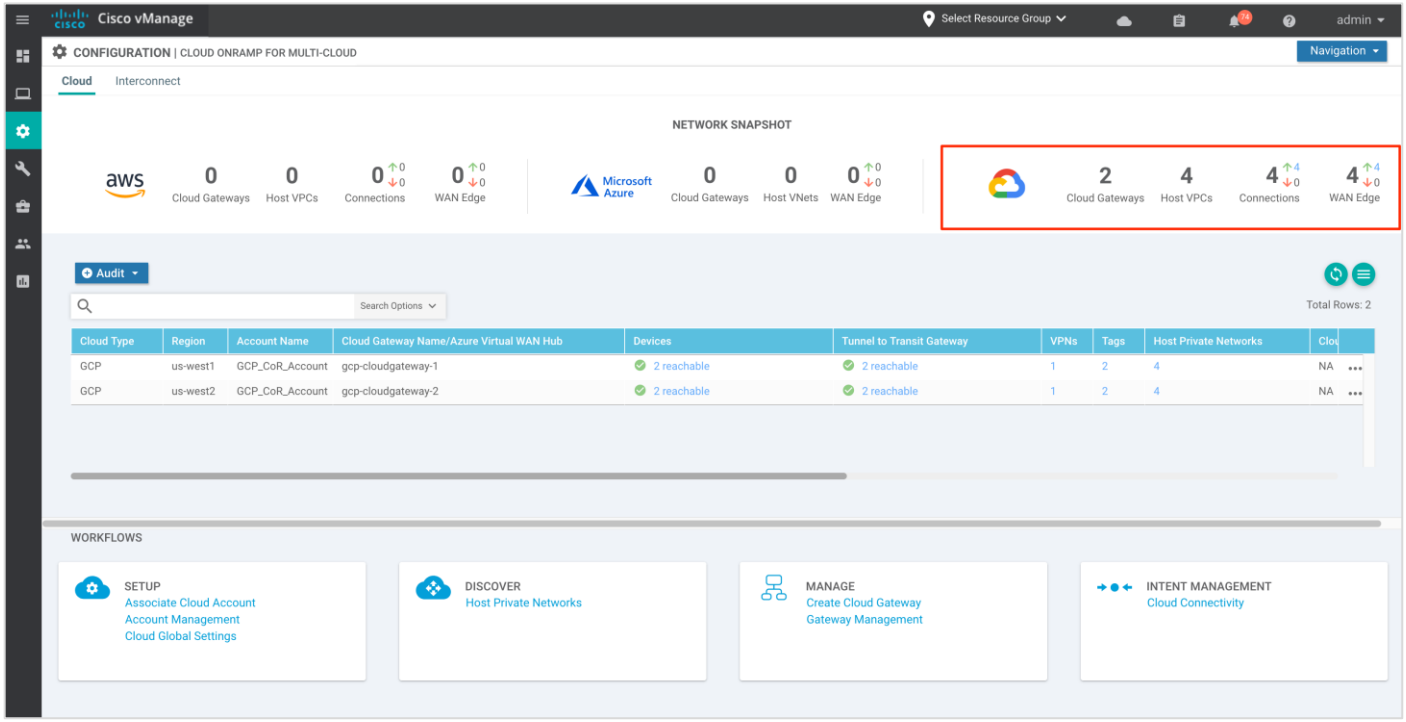

The upward arrow above the section **Connections** indicates the status of the Site-to-Cloud connectivity i.e. the mapping between the site-to-cloud VPN/ VRF to Google host VPC. Click on the arrow to view the number of

© 2020 Cisco and/or its affiliates. All rights reserved.

cloud gateways brought up, along with the VPNs label used to map cloud hosted VPC resources to the service side VPN in the Catalyst 8000v device.

In this design, two cloud gateways are brought up, one in US-West1 and the other in US-West2. Each cloud gateway contains a pair of Catalyst 8000v devices and each device is assigned one service side VPN (VRF 1) for Site-to-Cloud mapping. All Google host VPCs are mapped to VPN 1 and the result of this mapping/connectivity is displayed in the table below. The outer IP address in column five indicate the IP address assigned to the Catalyst 8000v interfaces part of the site-to-cloud VPN/VRF (VRF 1).

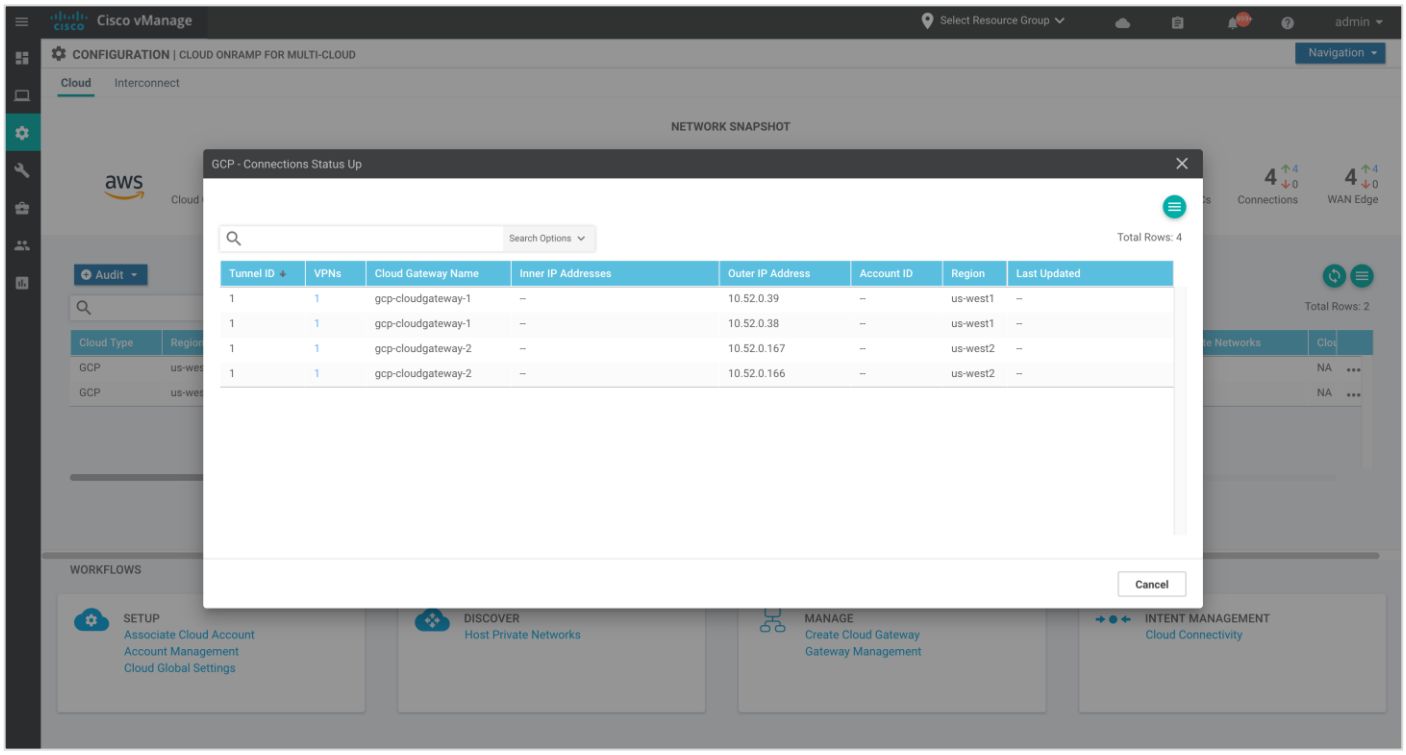

The upward arrow above the section WAN Edge indicates the reachable WAN Edge devices. Click the arrow to view additional details.

In this deployment, two cloud gateways are deployed each a pair of Catalyst 8000v devices of the following Chassis UUID. The table also lists out the System IP, the name given to each of the cloud gateways along with the Google region in which it is hosted.

Procedure 2: Deployed Google Cloud Resources

Step 1. Another way to monitor the resources deployed through the cloud onramp workflow is by scrolling across the section located at the middle of the Cloud onRamp for multi-cloud dashboard.

The columns list out the following

- Cloud Type: Lists out the type of cloud deployment, Azure, AWS or Google Cloud platform (GCP)
- Regions: Understand and validate the regions in which your Google cloud gateways are deployed
- Account Name: The name assigned while associating your google account to vManage NMS.
- Cloud Gateway Name: View the name of your cloud gateways deployed in each Google region.
- Device: This column contains a green tick symbol (success) or a red crossed out symbol (failure) to indicate the status of the reachable Catalyst 8000v devices in each of the two cloud gateways.

© 2020 Cisco and/or its affiliates. All rights reserved.

- Tunnel to Transit Gateway: This column indicates the connectivity status between your site-to-cloud service side VPN in Catalyst 8000v to the host VPC mapping.
- VPNs Tags and Host Private Networks: These three columns indicate the VPNs to tag mapping, along with details regarding the host VPCs discovered and attached to each tag.

Step 2. To gather further details regarding the Google cloud resources deployed through the Cloud onRamp workflow, click on the three dots **(...)** located to the right side.

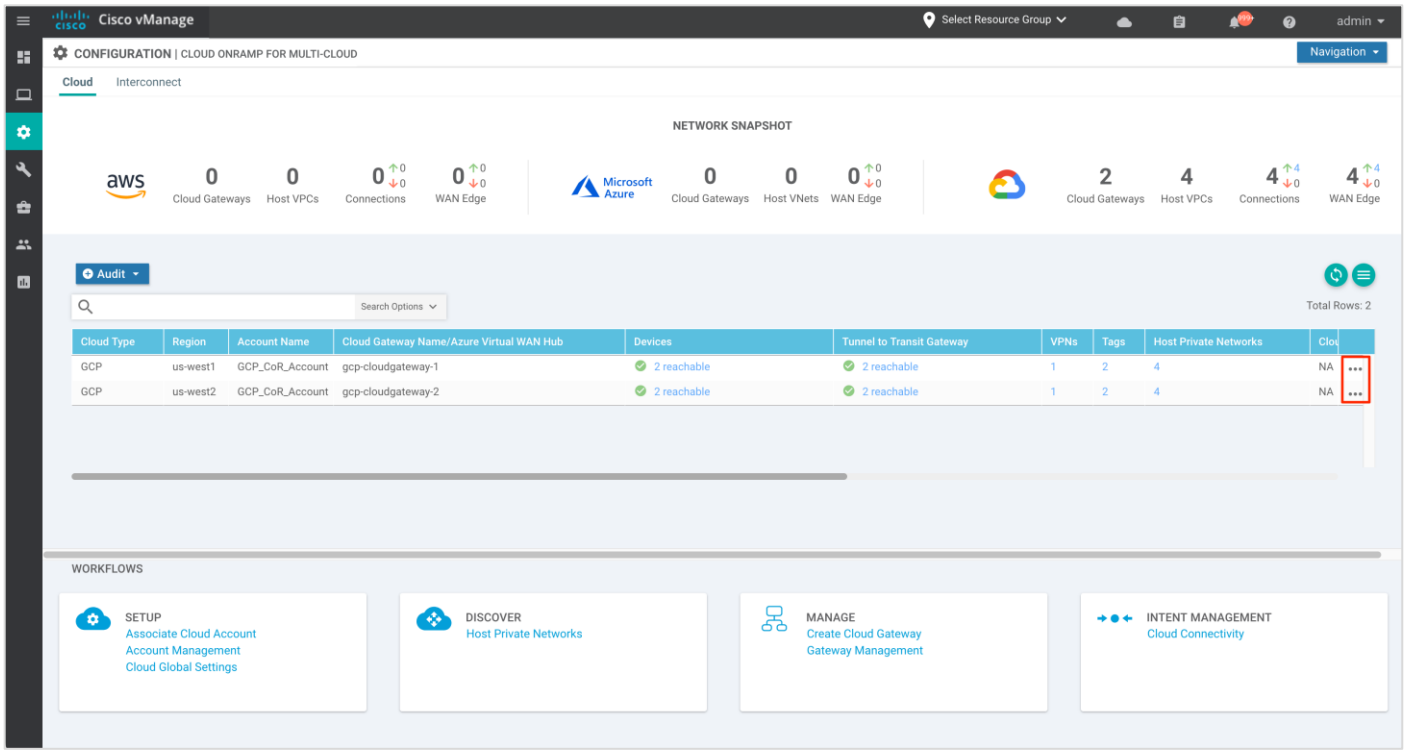

Step 3. Select **Intent Realization Summary** from the drop-down to view the mapped and unmapped host VPCs.

© 2020 Cisco and/or its affiliates. All rights reserved.

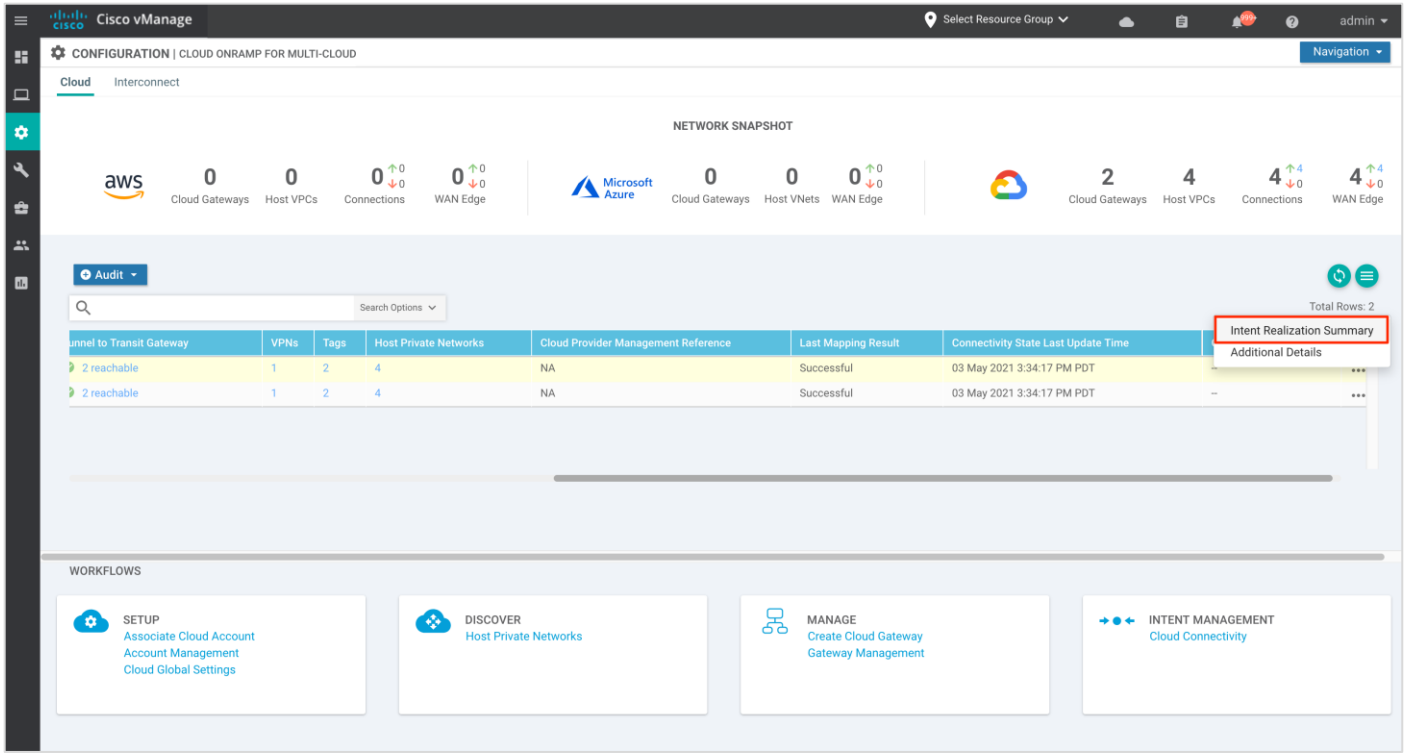

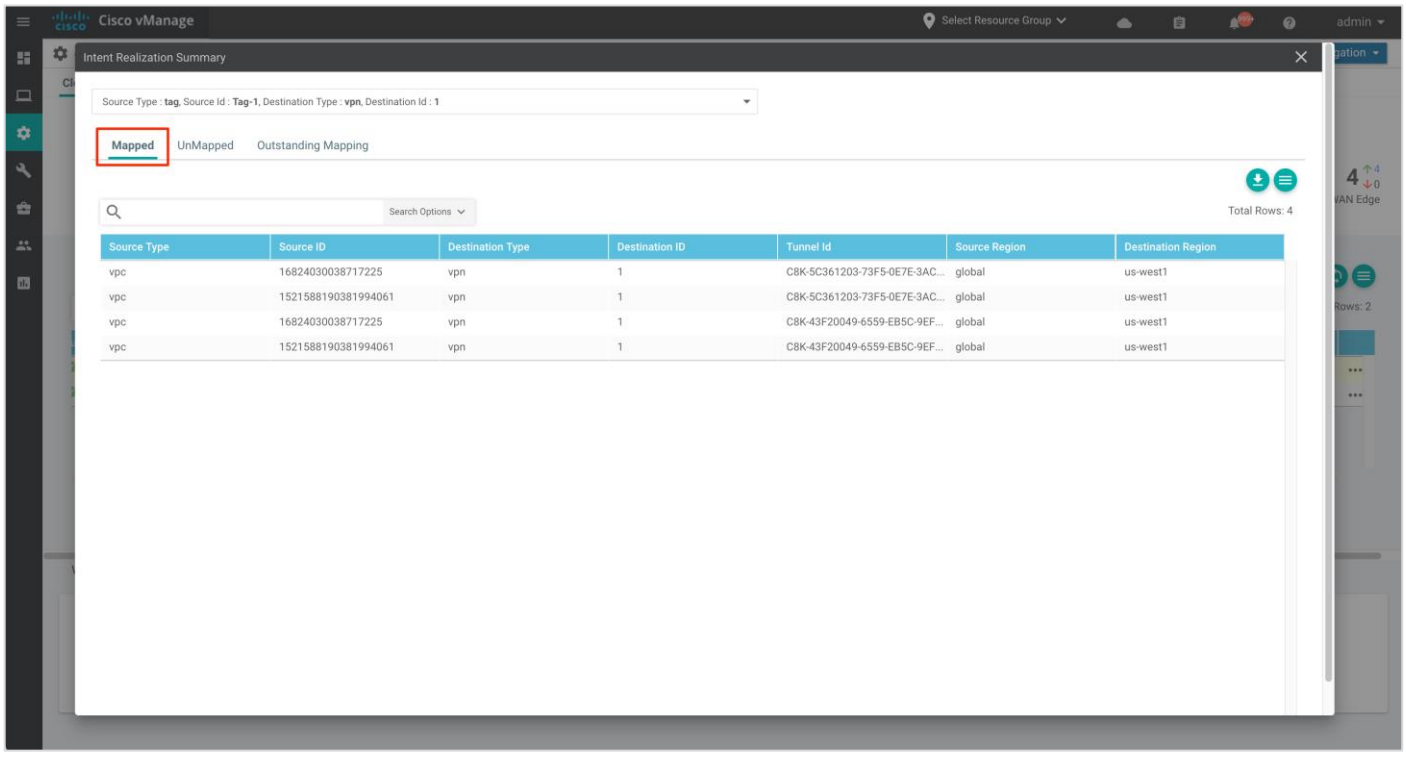

Step 4. To view Google cloud resources deployed through this workflow, select **Additional Details**.

© 2020 Cisco and/or its affiliates. All rights reserved.

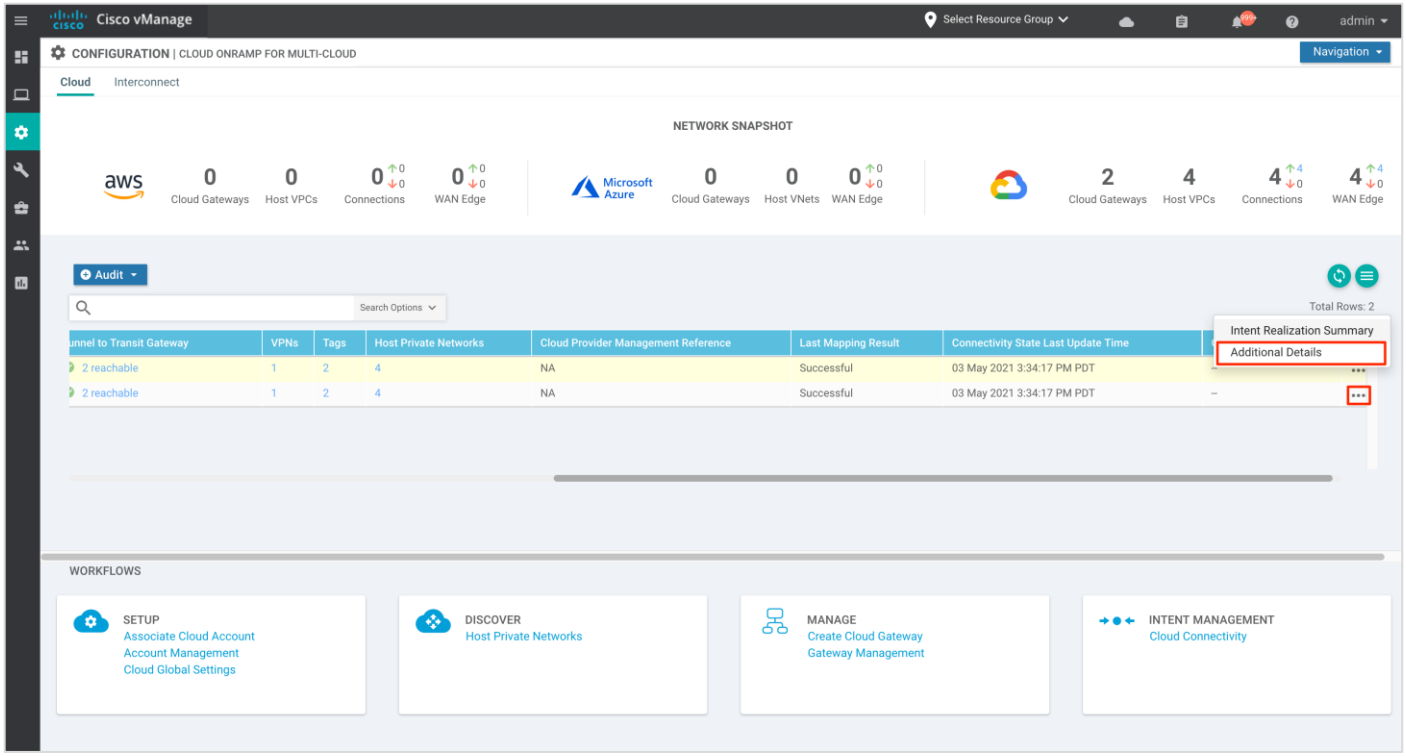

The pop-up window lists out all the google resources deployed through the site to cloud and site to site VPC workflow.

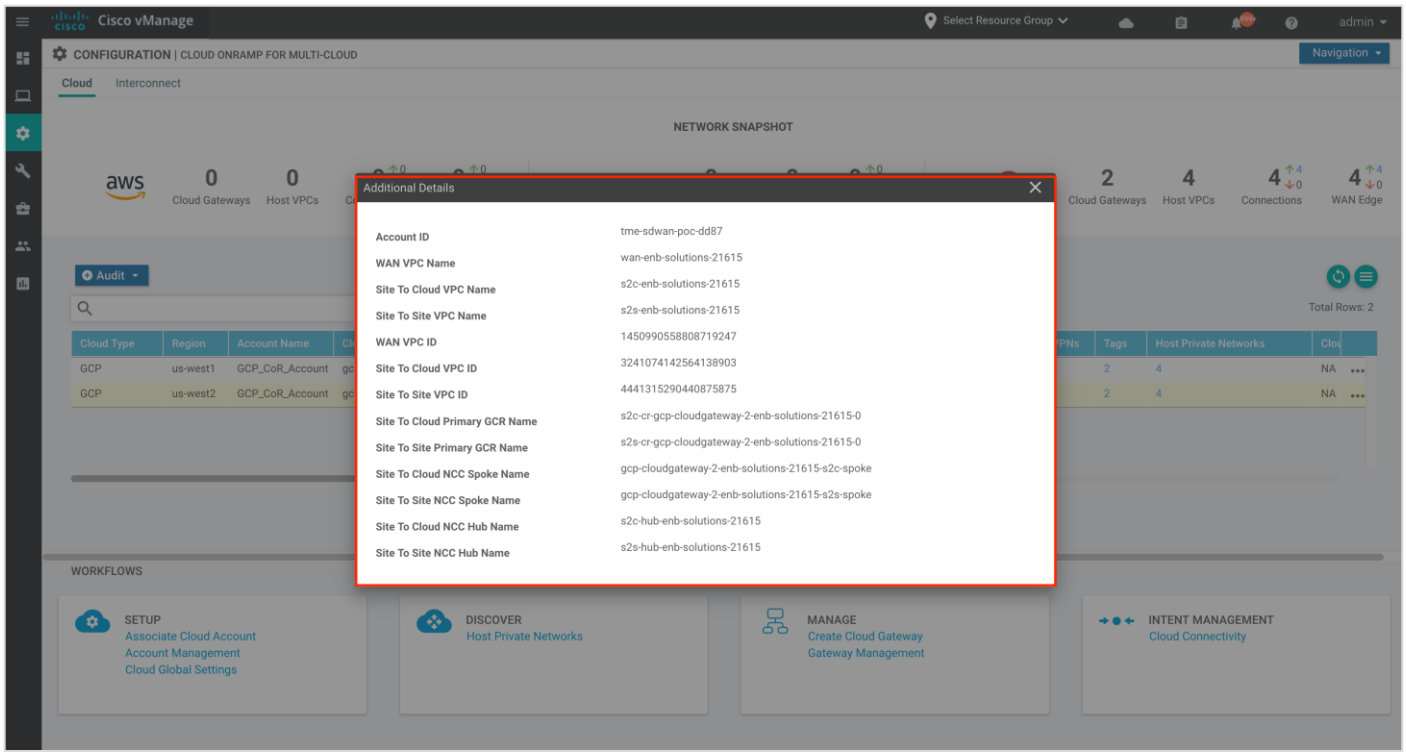

© 2020 Cisco and/or its affiliates. All rights reserved.

Procedure 3: Manage Google Billing ID and Credentials

Step 1. To view, manage and update the Billing ID and credentials, click **Account Management** under the section **Setup.**

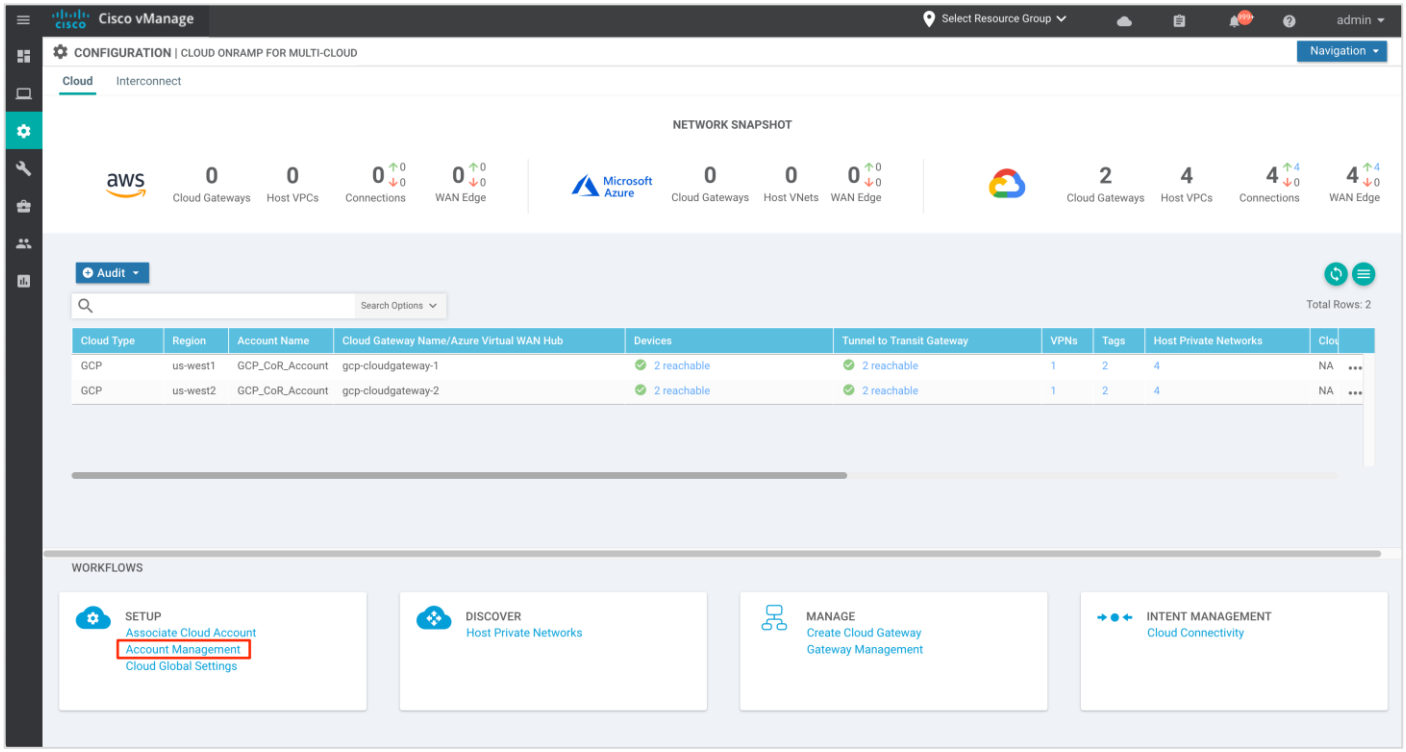

Step 2. Under the section **Cloud**, the Google cloud account details associated with your vManage NMS are listed. To manage the cloud account details, click on the three dots **(...)** located on the right corner.

Select:

- View to manage all the cloud account details.
- Update Cloud Account Details to edit the Billing ID details.
- Update Cloud Account Credentials to update the credentials associated with vManage.
- Remove to delete the Account Details.

© 2020 Cisco and/or its affiliates. All rights reserved.

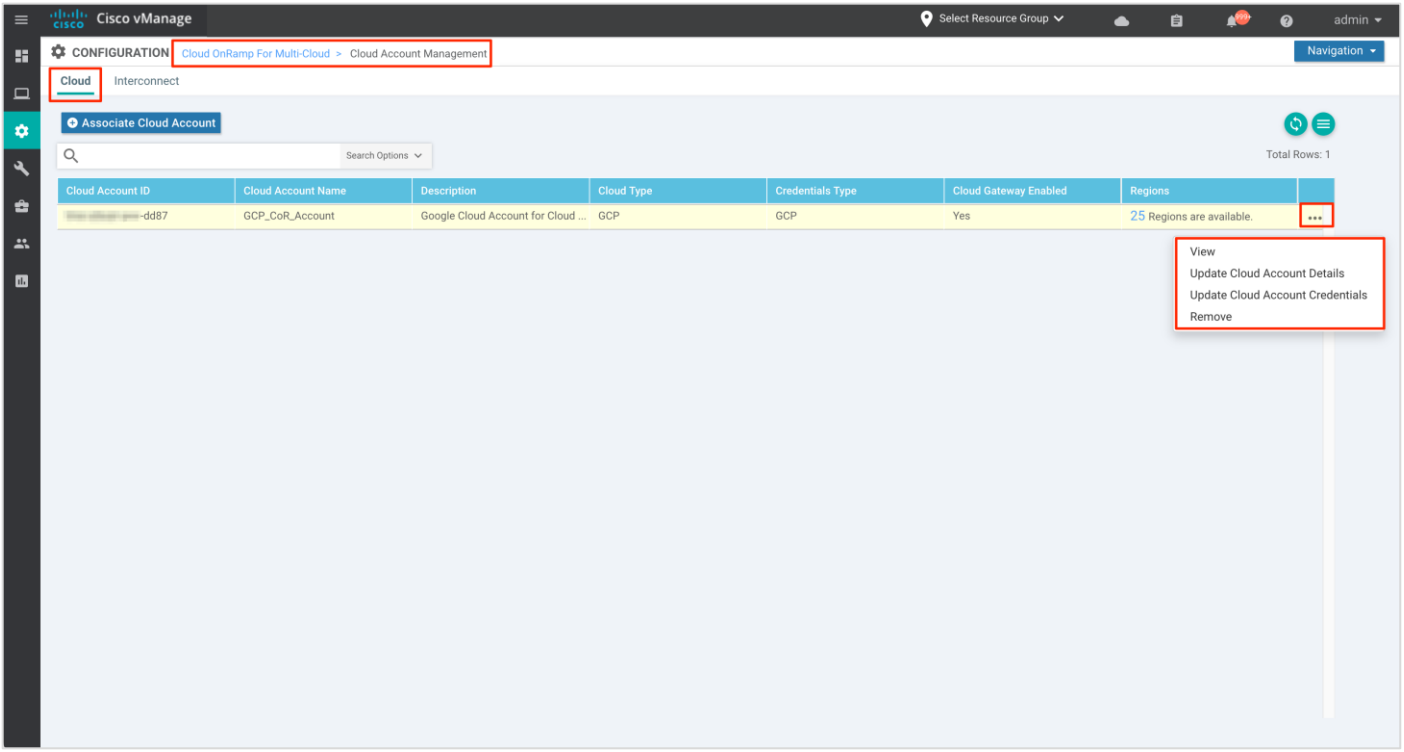

The example screenshot displays all the entered Cloud Account Details, located under **Cloud onRamp for Multi-Cloud** > **Cloud Account Management** > **View Cloud Account**.

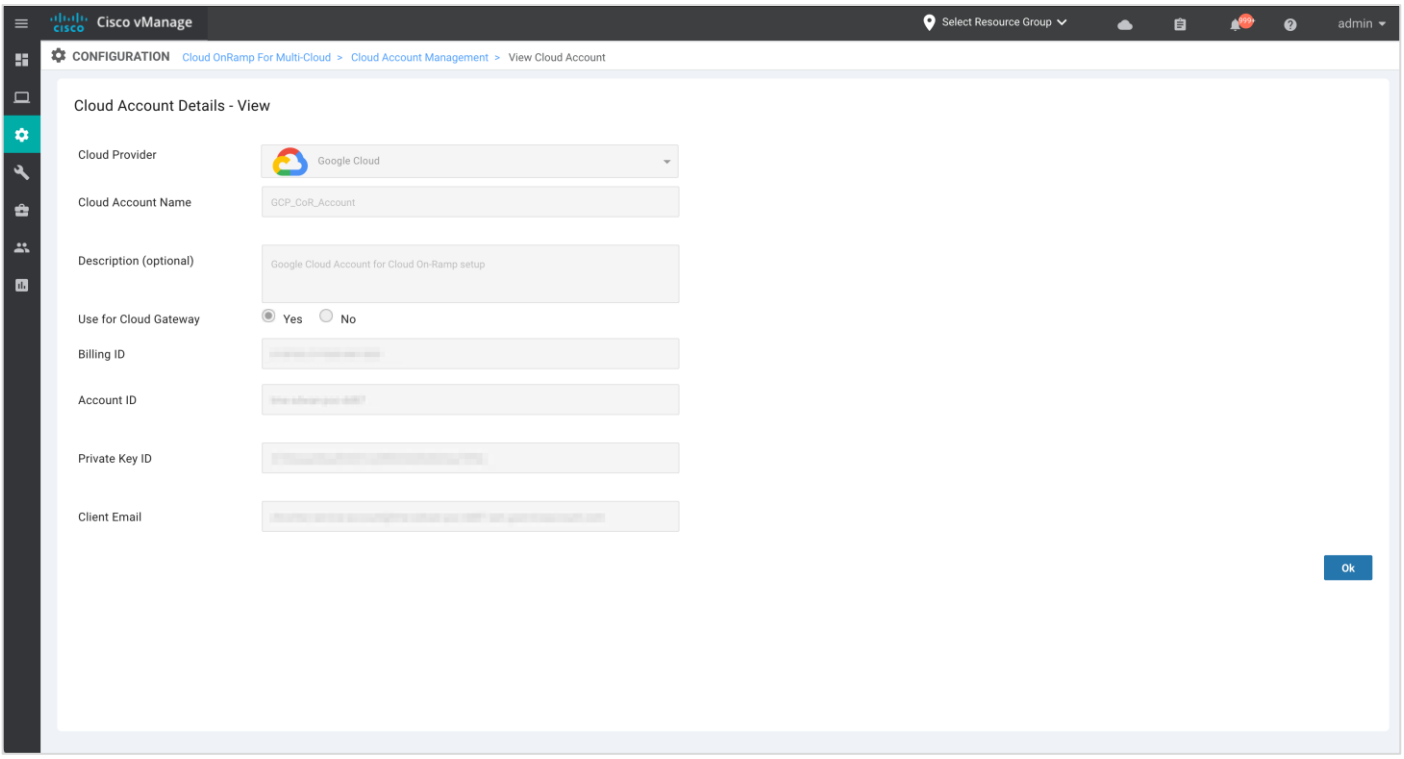

Procedure 3: Manage Cloud Global Settings

Step 1. To view or delete the Global Cloud Gateway settings. Under **Manage**, click on **Gateway Management**.

© 2020 Cisco and/or its affiliates. All rights reserved.

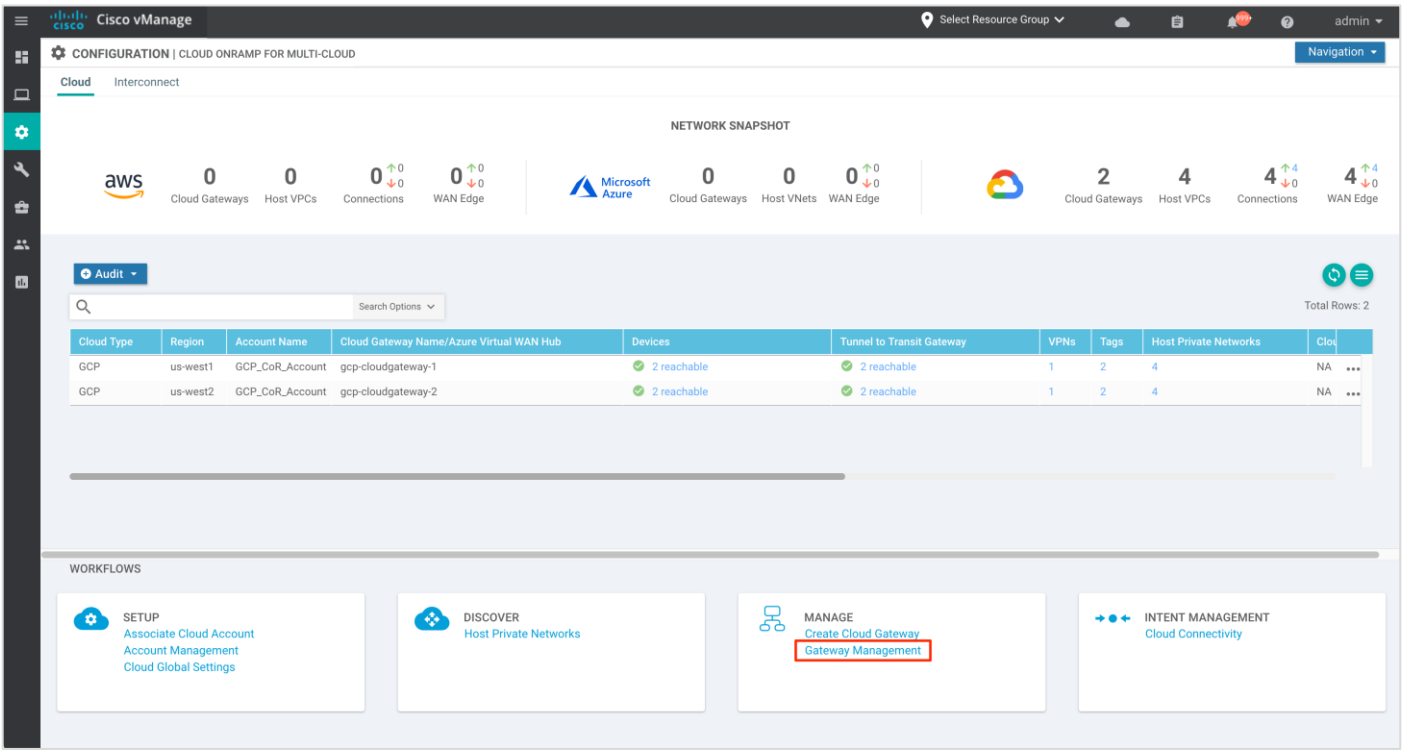

The example screenshot below displays two cloud gateways created through the workflow.

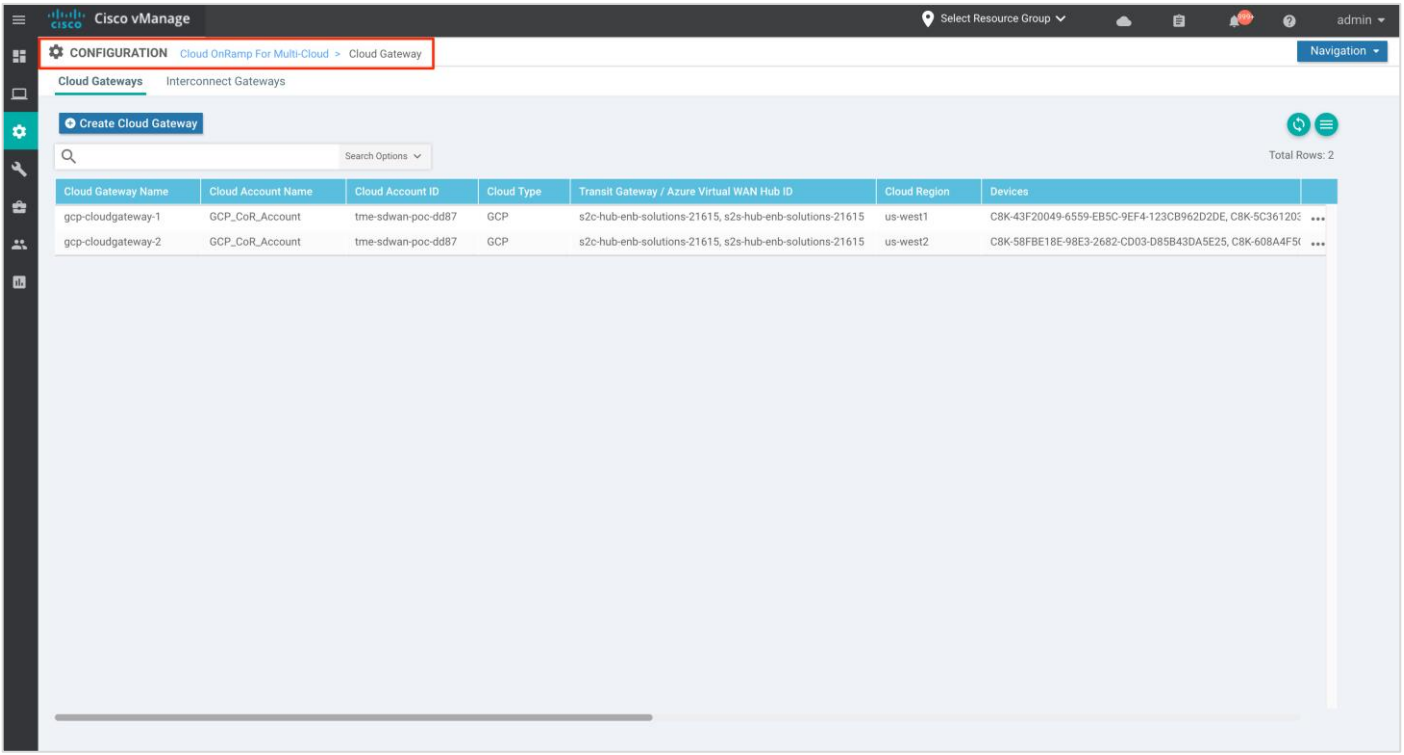

Click on the three dots **(...)** on the right to **View** or **Delete** the cloud gateways.

© 2020 Cisco and/or its affiliates. All rights reserved.

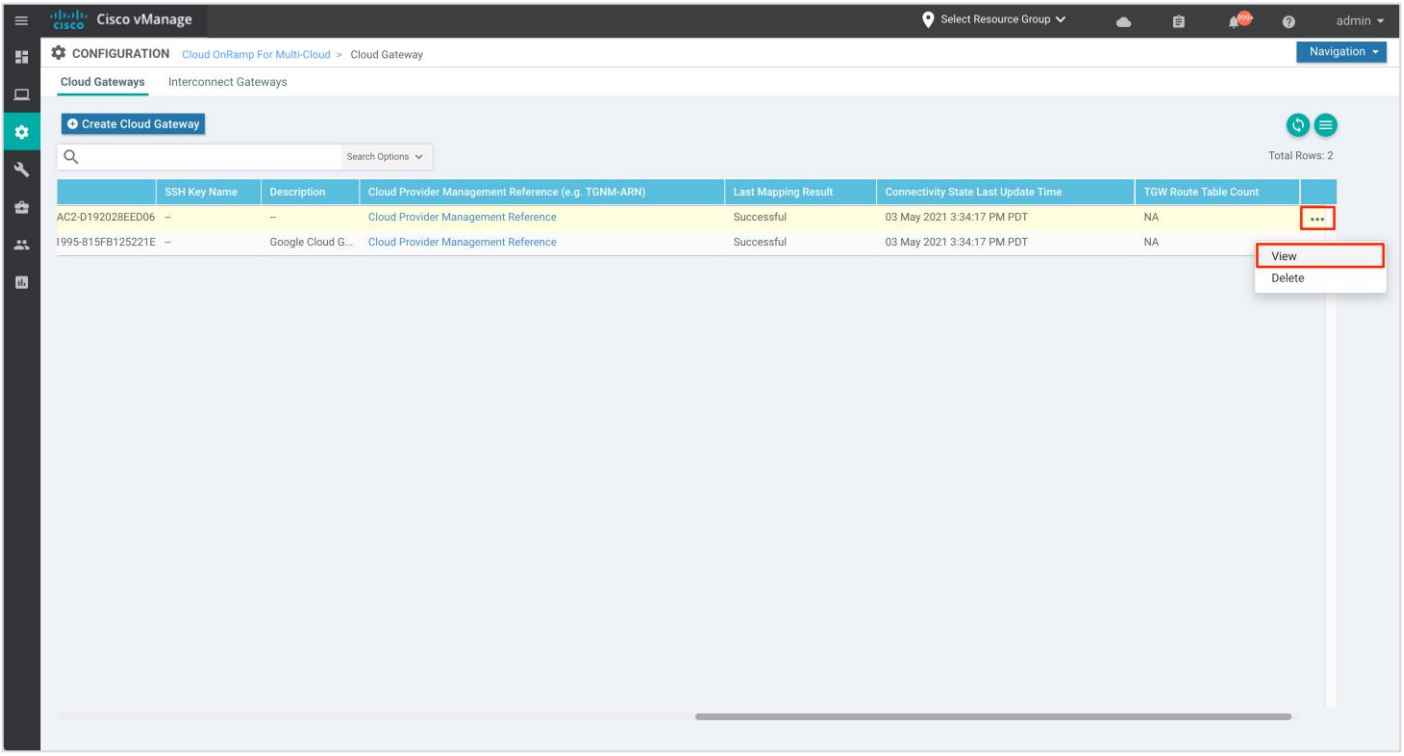

The example screenshot below displays the details entered within the cloud gateway.

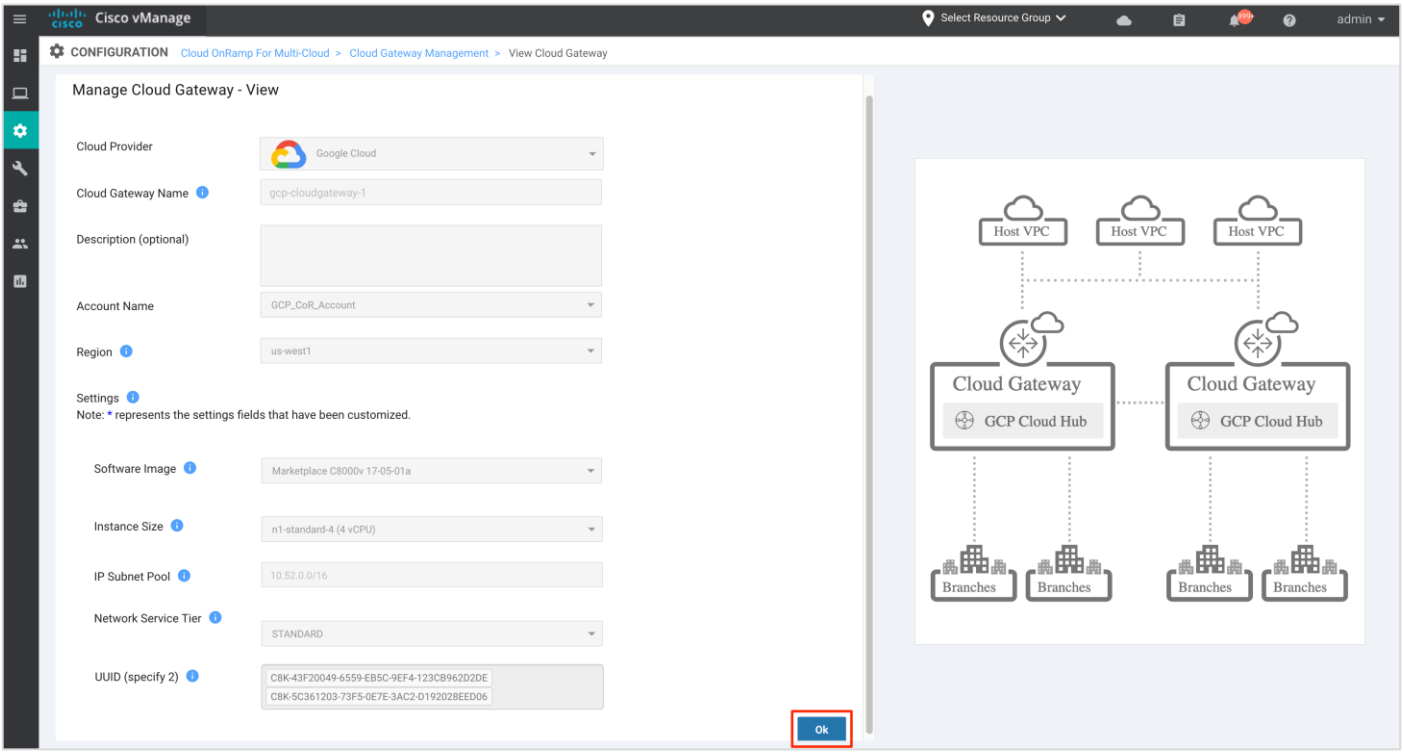

Procedure 3: Manage VPN to Tag Mapping

To edit or monitor an existing mapping between your service side VPN and a tag follow the steps below

© 2020 Cisco and/or its affiliates. All rights reserved.

Step 1. Under **Intent Management**, click **Cloud Connectivity** to view and/ or edit the mapping of the host VPCs in Tags to VPN.

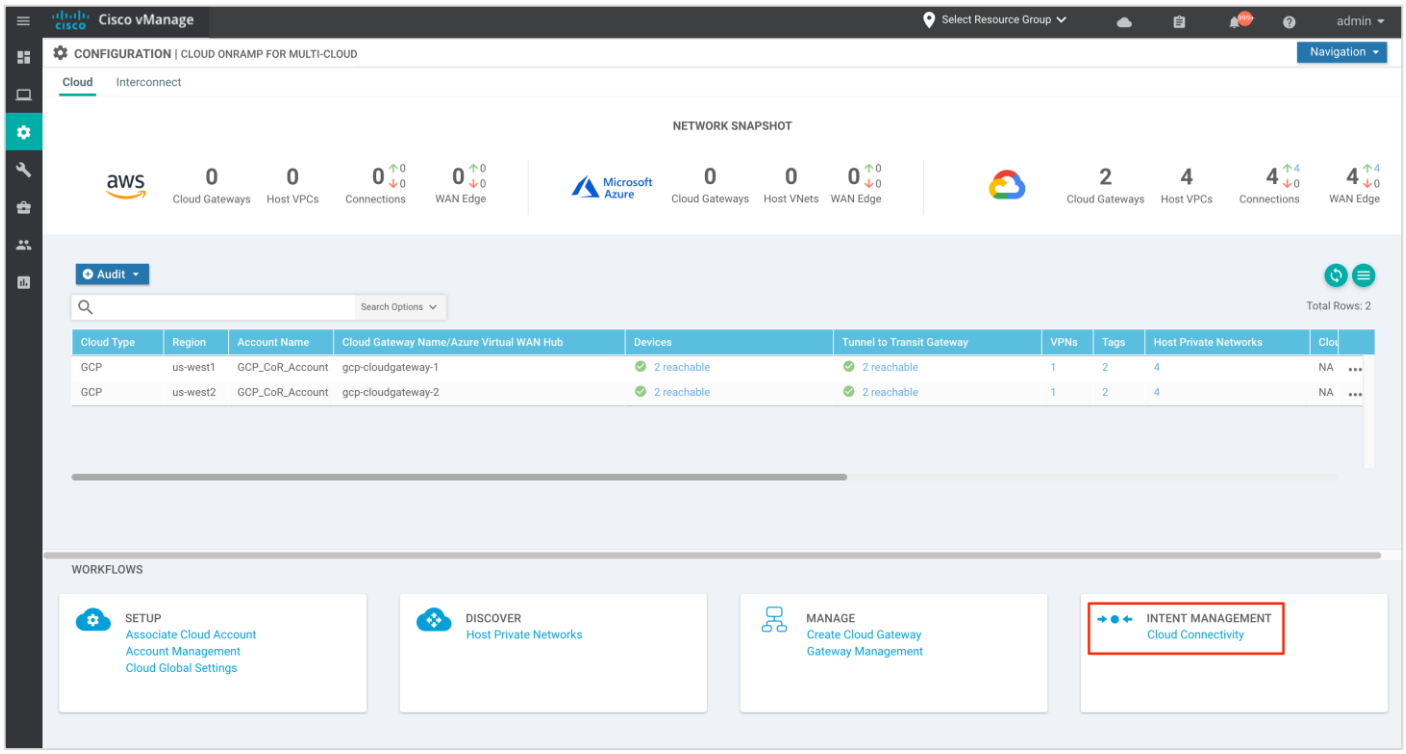

Step 2. To associate additional tags that contain host VPCs to VPN, click **Edit** located on the right.

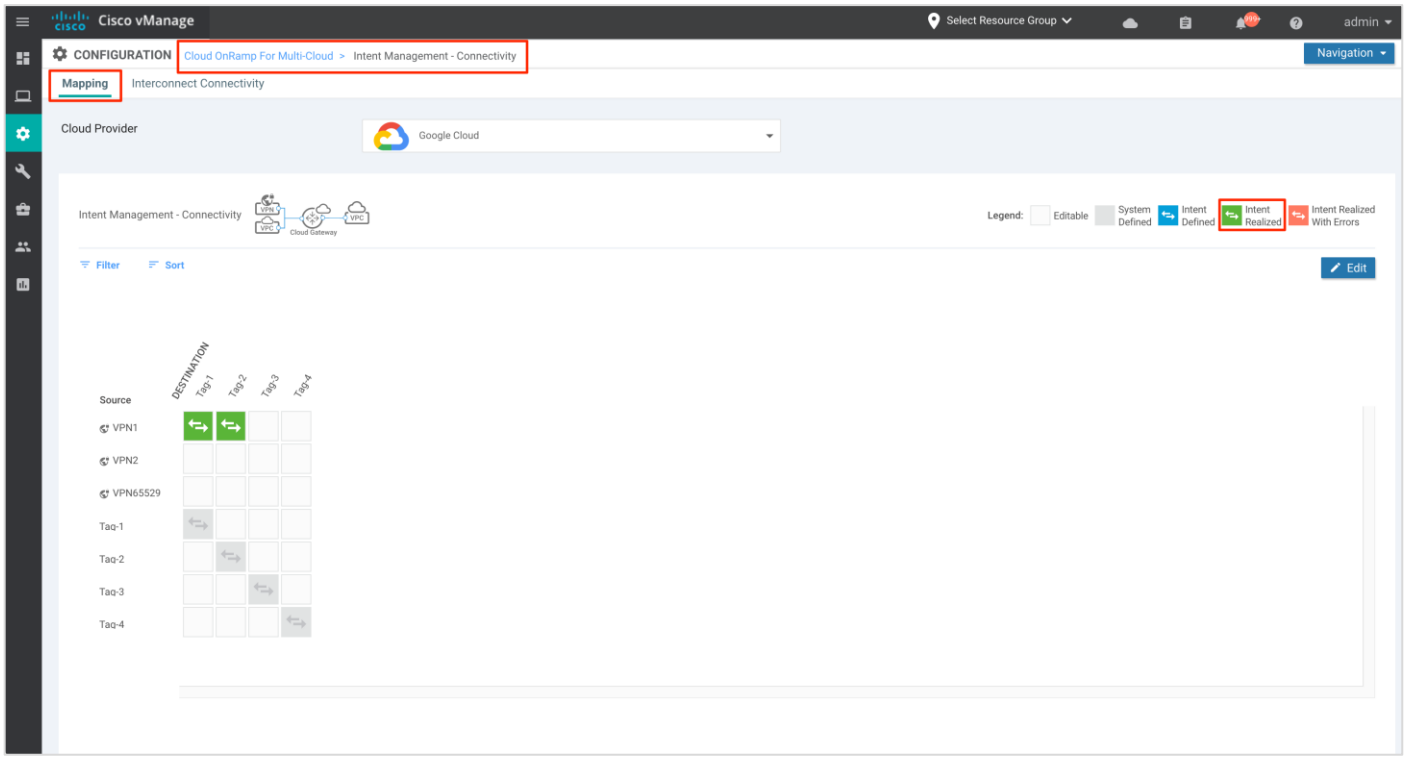

Step 3. Click **Filter** to view the mapping based on a certain source VPN/ Tag and destination Tag.

© 2020 Cisco and/or its affiliates. All rights reserved.

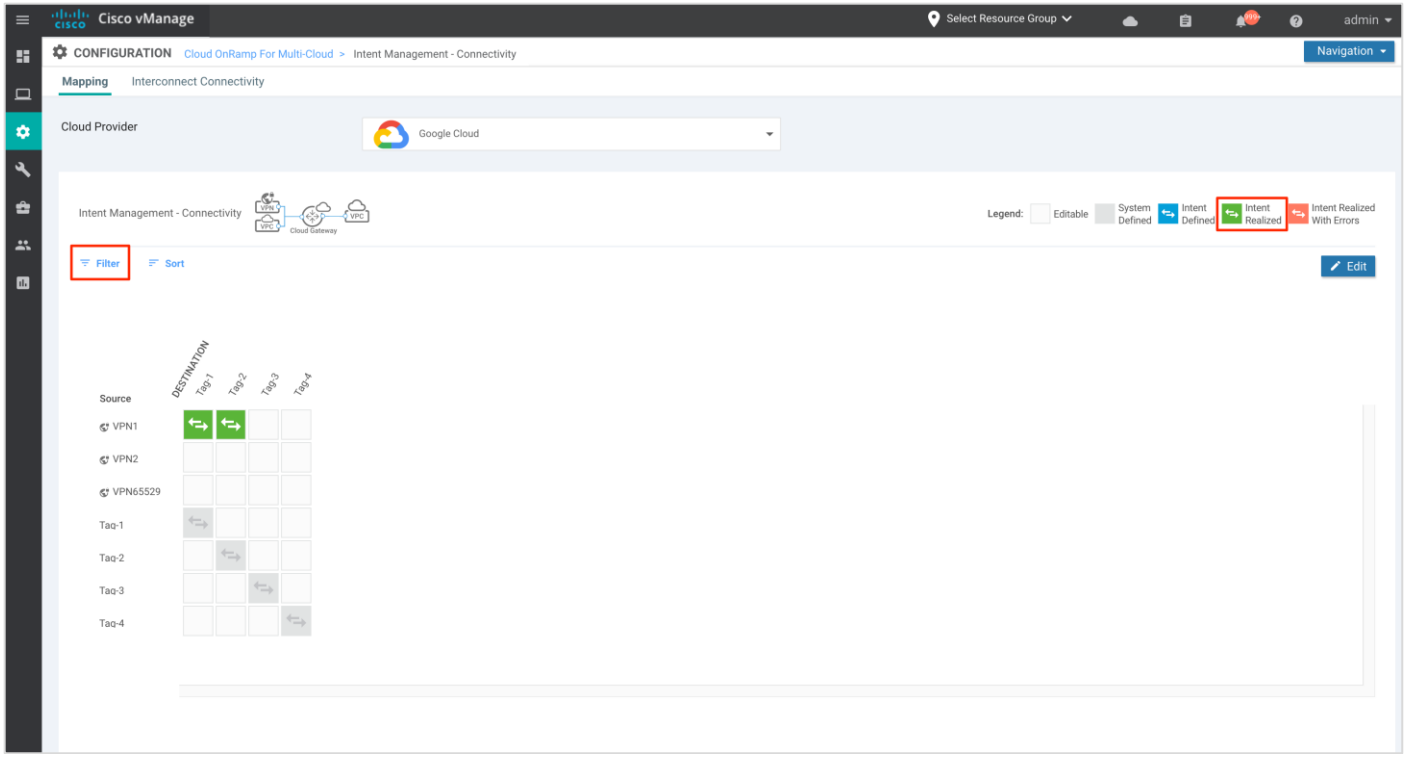

The example screenshots below display the mapping for source VPN 1 and Tag1 mapping.

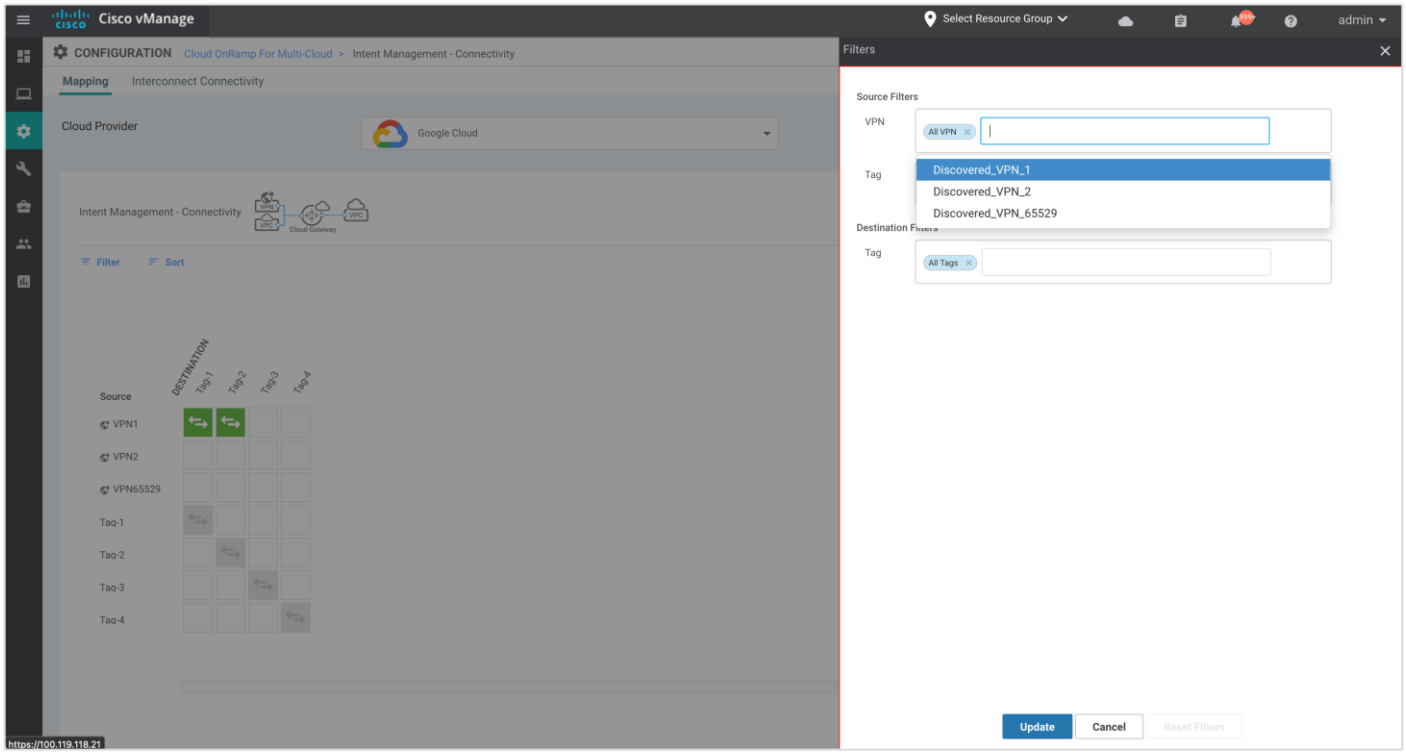

© 2020 Cisco and/or its affiliates. All rights reserved.

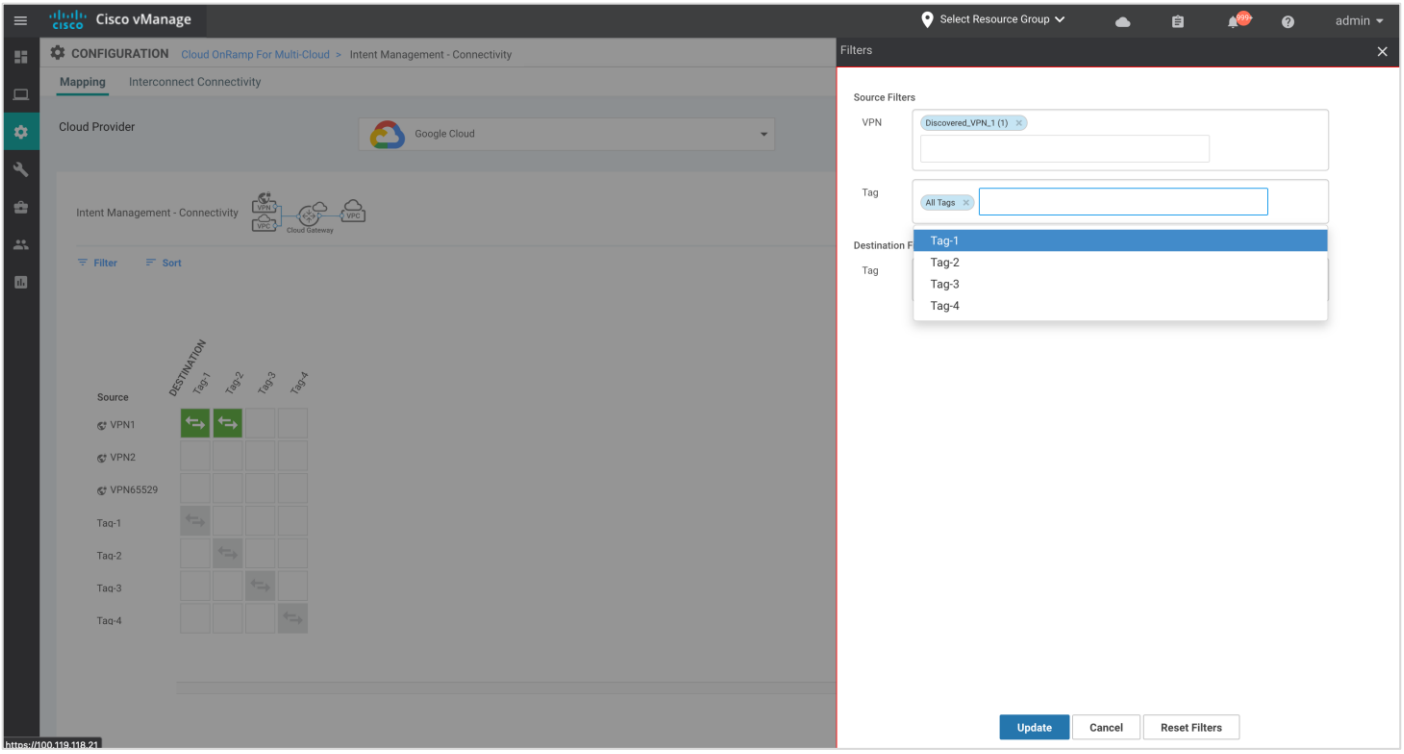

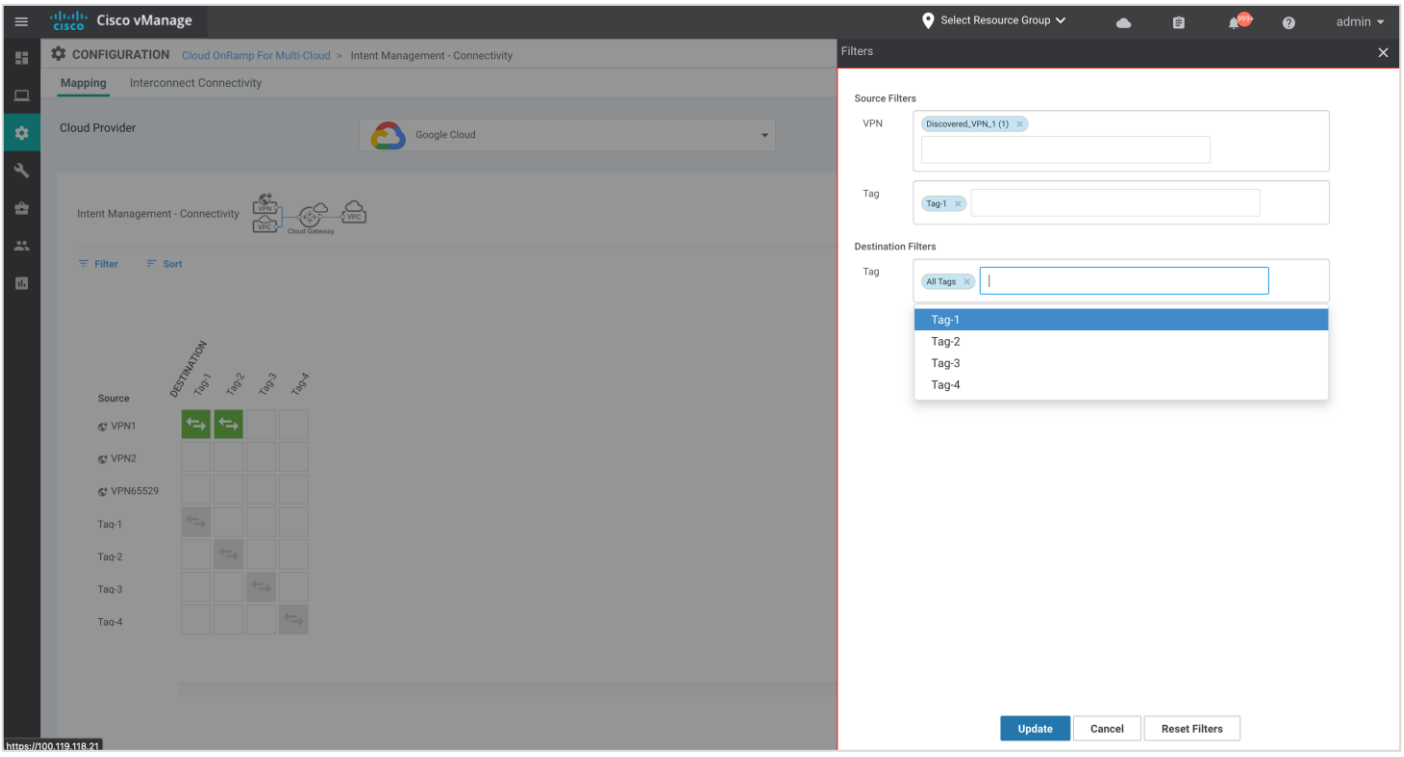

© 2020 Cisco and/or its affiliates. All rights reserved.

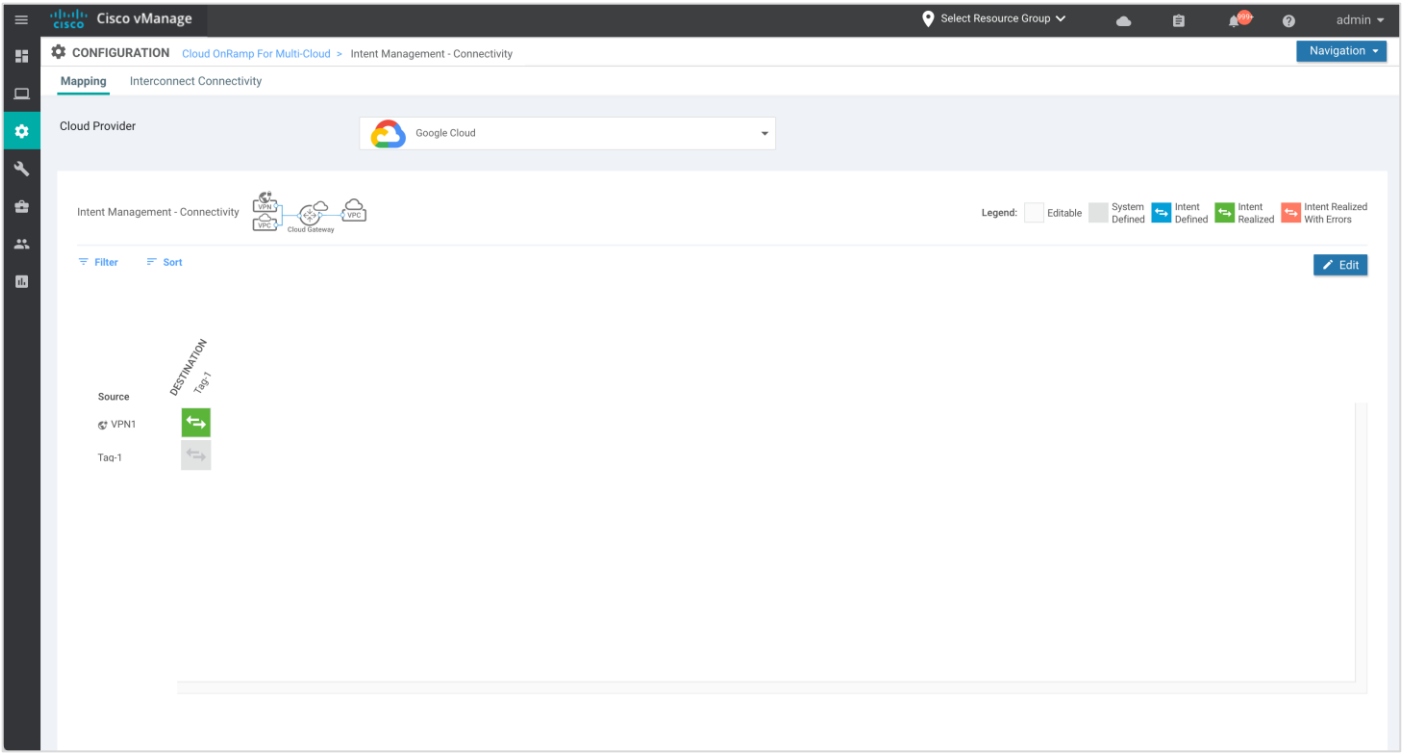

Step 4. You can also click on **Sort** to further filter/ arrange the view.

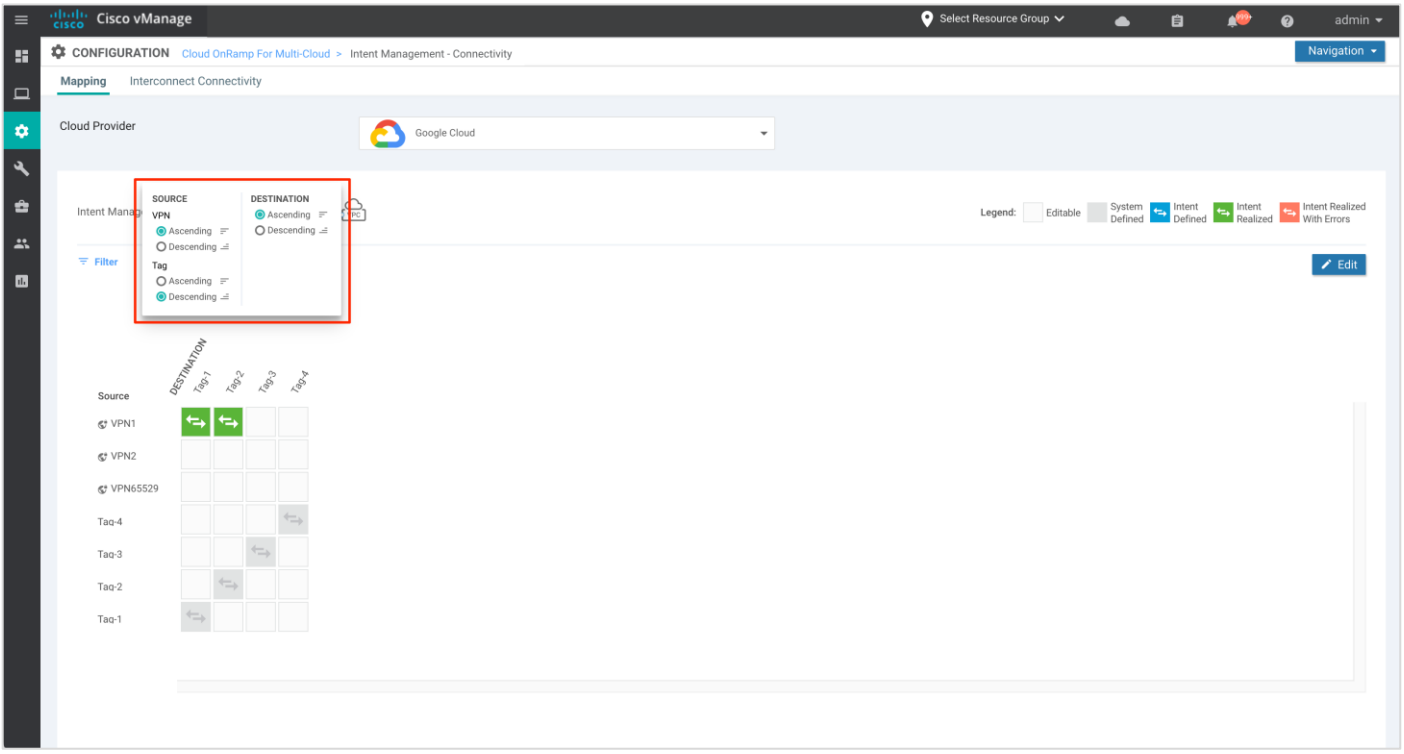

© 2020 Cisco and/or its affiliates. All rights reserved.

Process 2: vManage Monitor Dashboard

Procedure 1: Monitor SD-WAN devices via Monitor tab

Step 1. To monitor both your controllers and WAN Edge devices within the Transit VPC, Navigate to **Monitor** > **Real Time**. The **Device Options** tab lists outputs like logs gathered from the show output.

Some example outputs are added below:

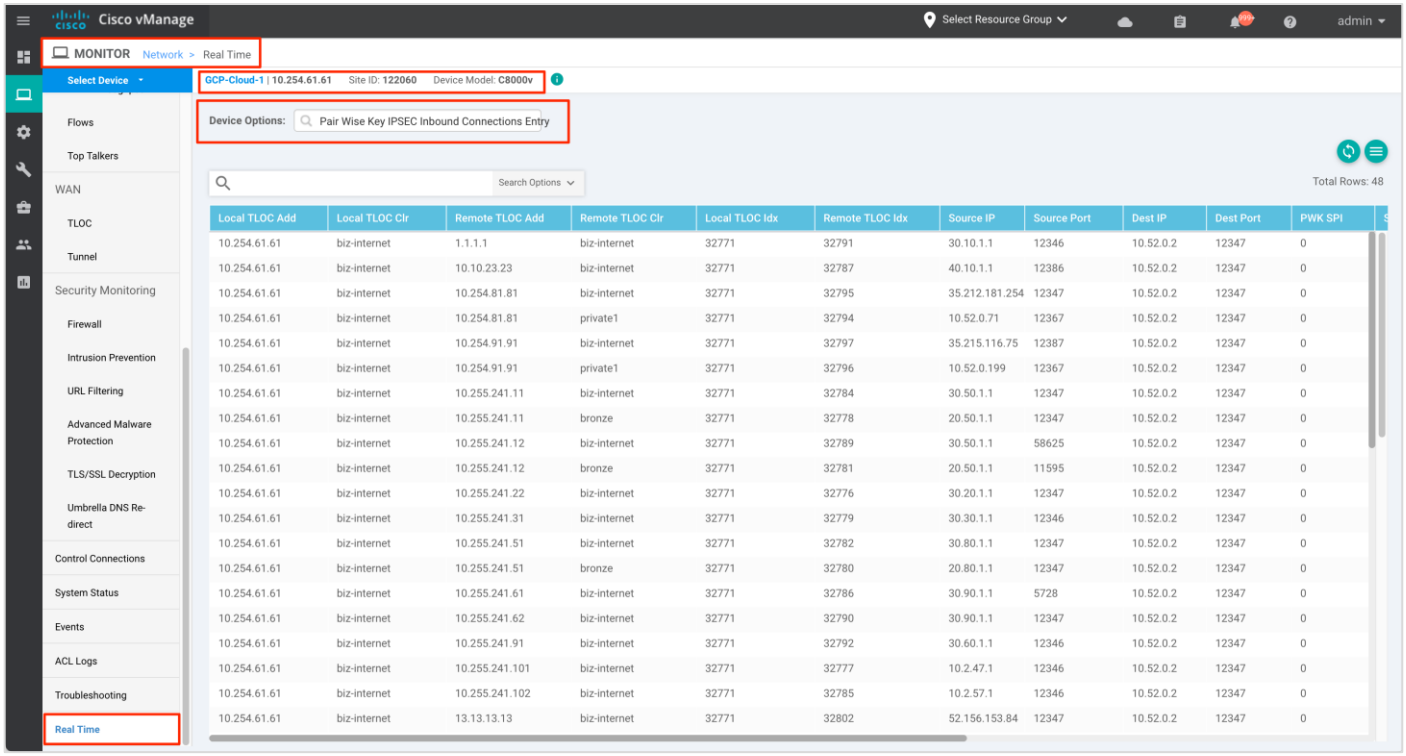

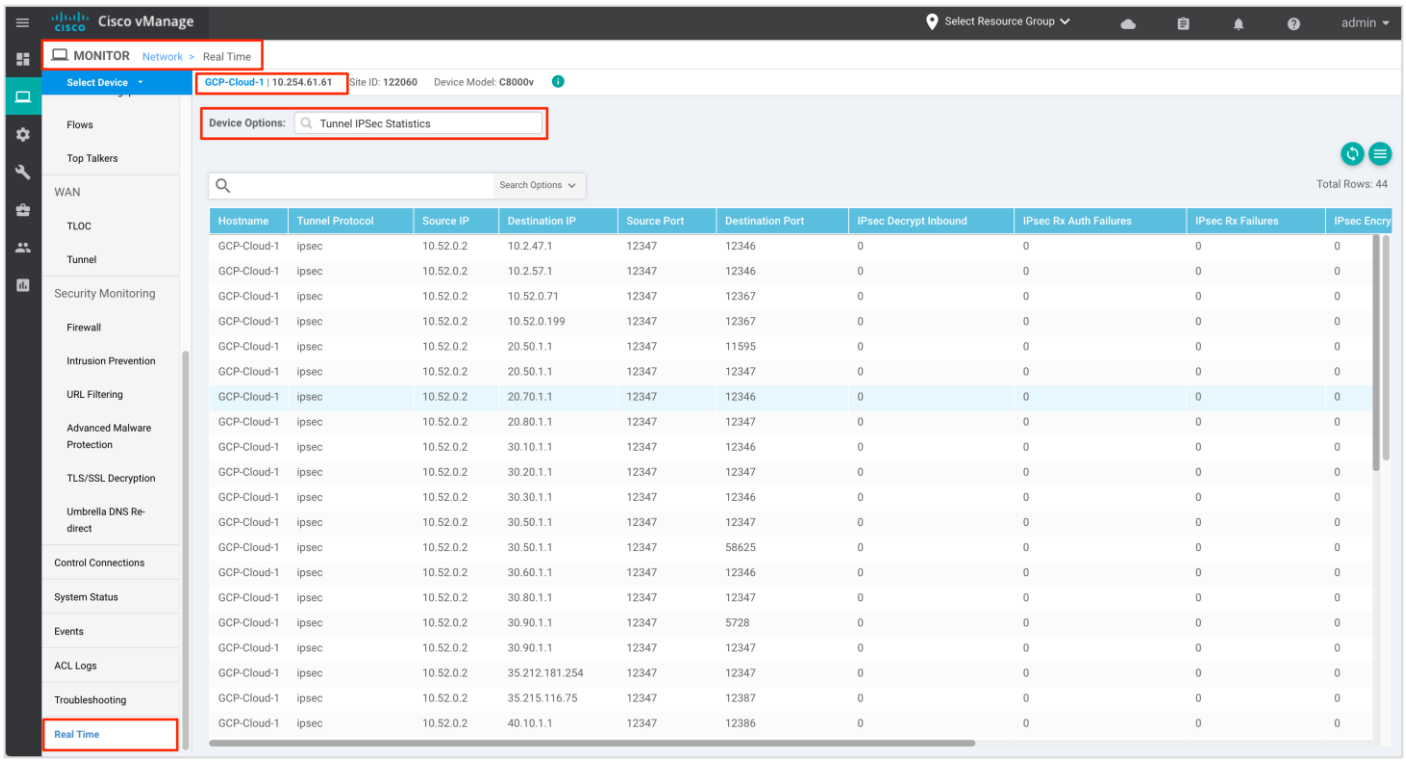

© 2020 Cisco and/or its affiliates. All rights reserved.

Step 2. Navigate to **Monitor** > **Tunnel**, to monitor both your controllers and WAN Edge devices within the Transit VPC.

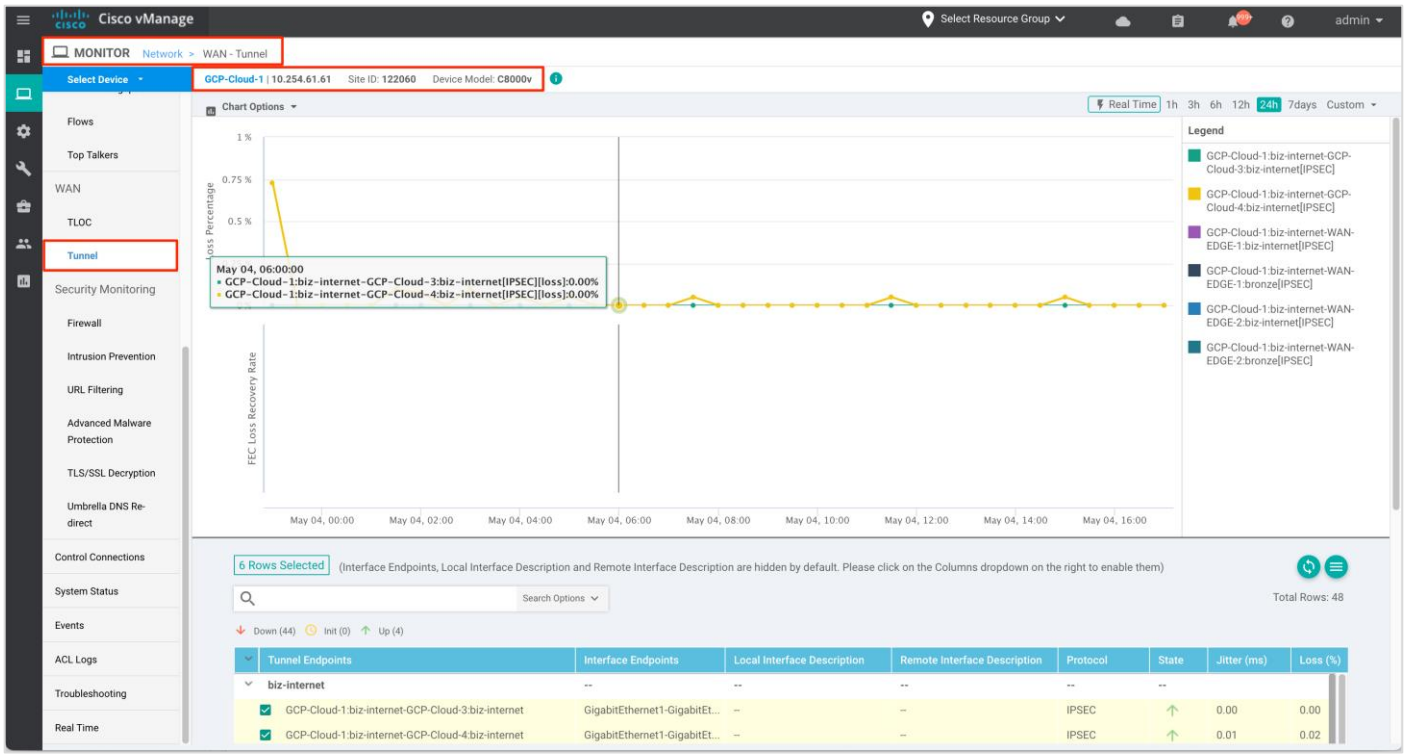

Step 3. Navigate to **Monitor** > **Events**, to monitor both your controllers and WAN Edge devices within the Transit VPC.

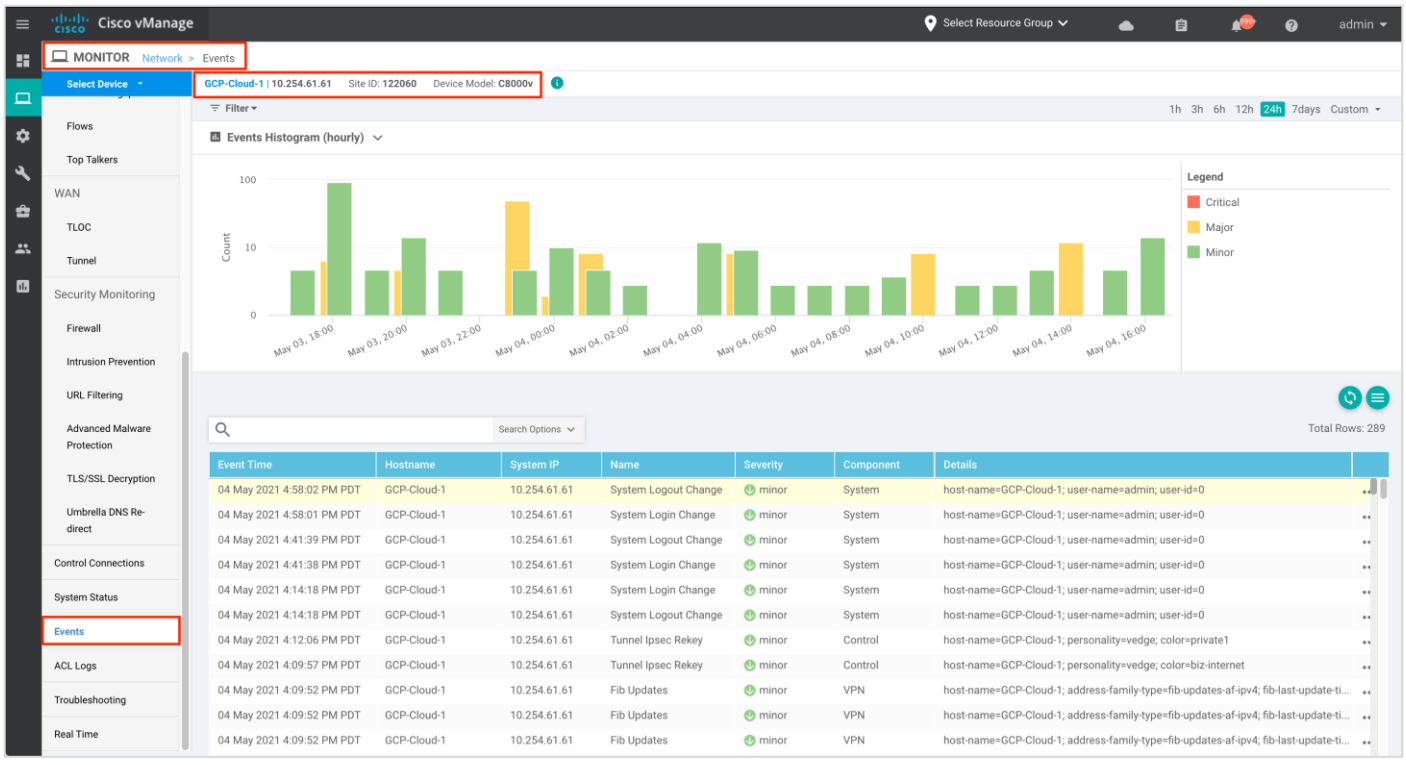

© 2020 Cisco and/or its affiliates. All rights reserved.

Process 3: vManage SSH Dashboard

Procedure 1: Monitor Catalyst 8000v devices via CLI

Step 1. To view the IP address and BGP peers in the global routing table, issue the following commands “show ip int br” and “show bgp summary”.

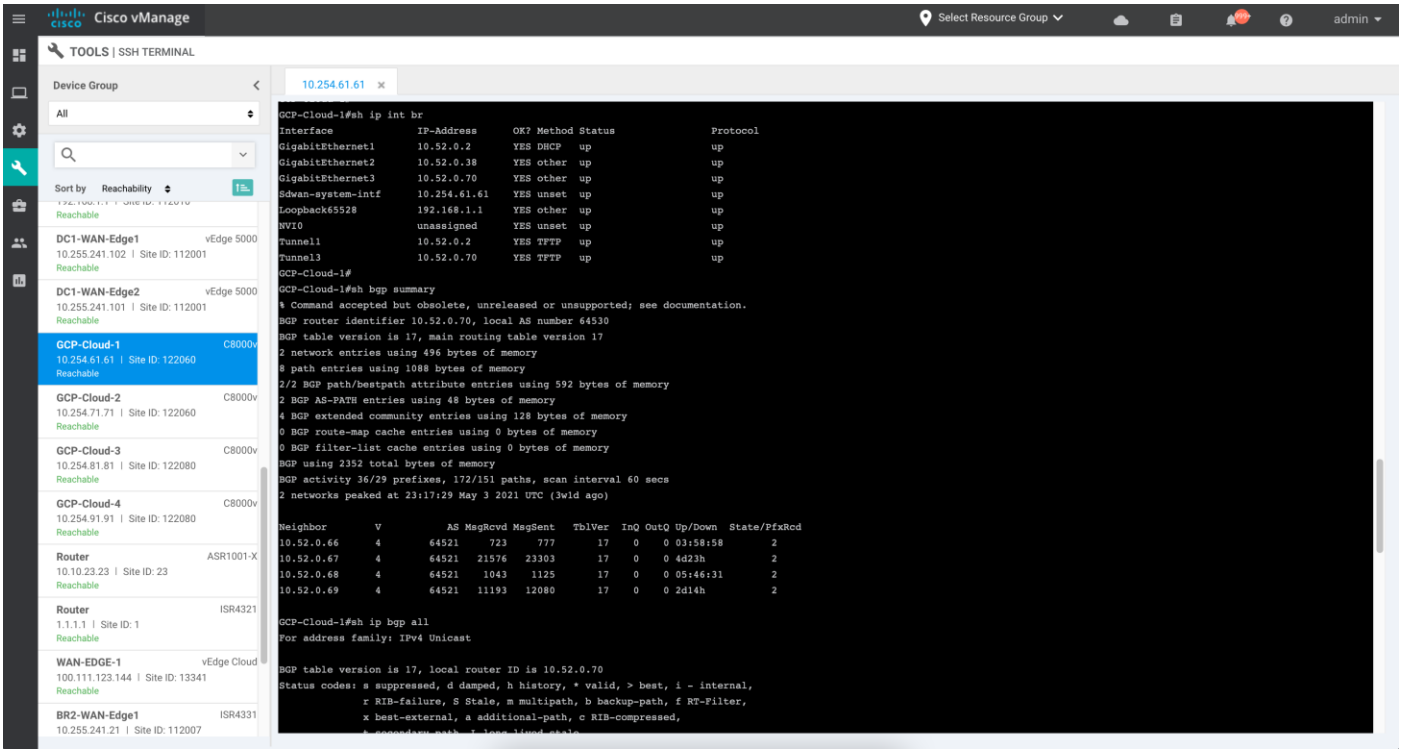

Step 2. To view the BGP routes learnt, enter CLI “show ip bgp all”.

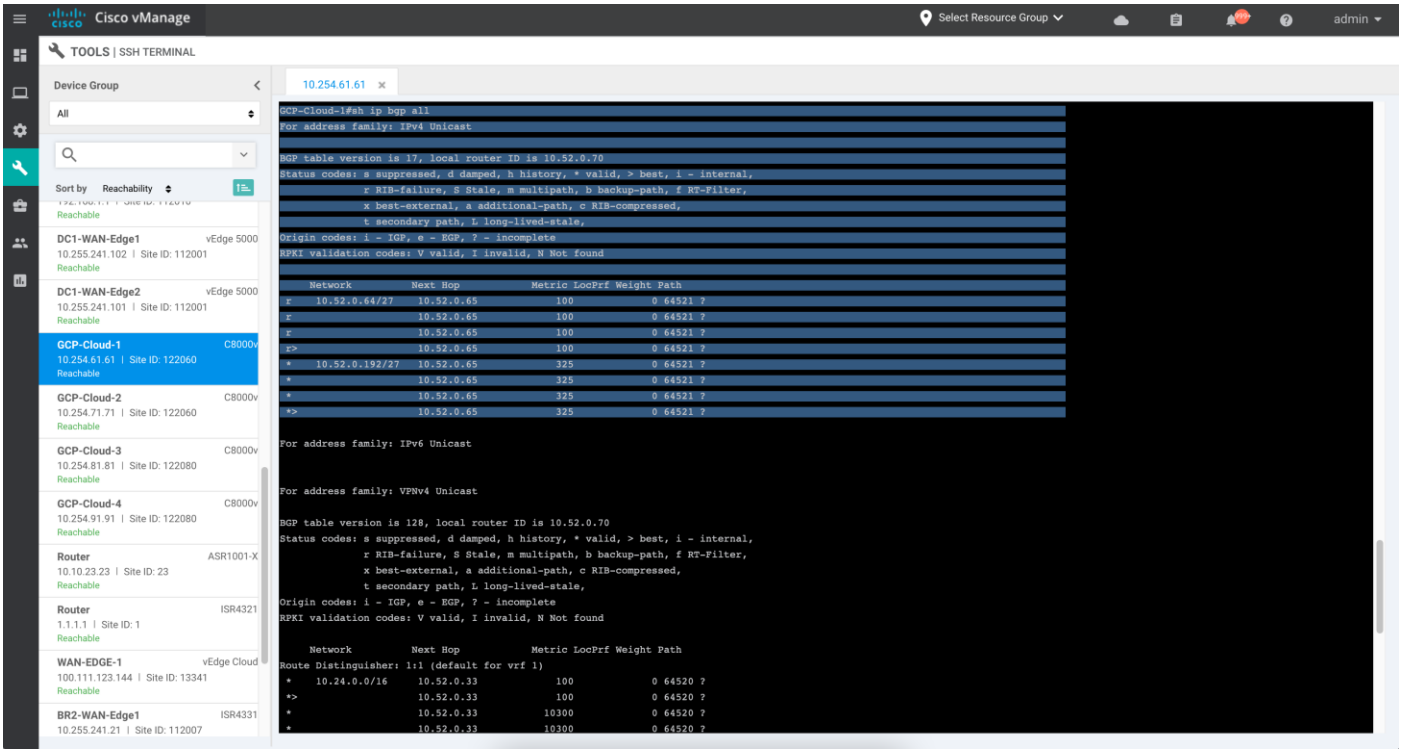

Step 3. To view the BGP routes learnt in VRF 1, issue the following commands “show ip bgp vpvnv4 vrf 1”.

© 2020 Cisco and/or its affiliates. All rights reserved.

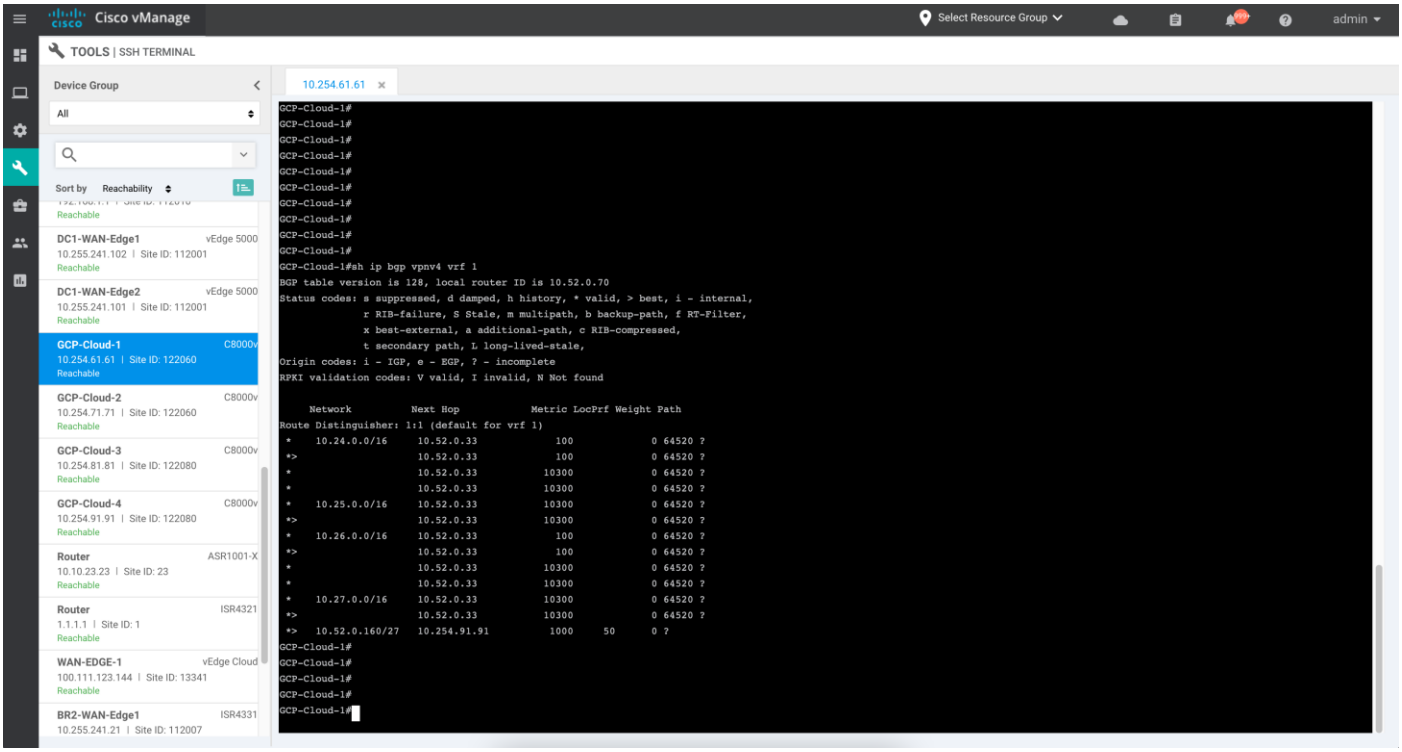

Step 4. To view BGP neighbors learnt within service side, enter CLI “show ip bgp vpvnv4 all summary”.

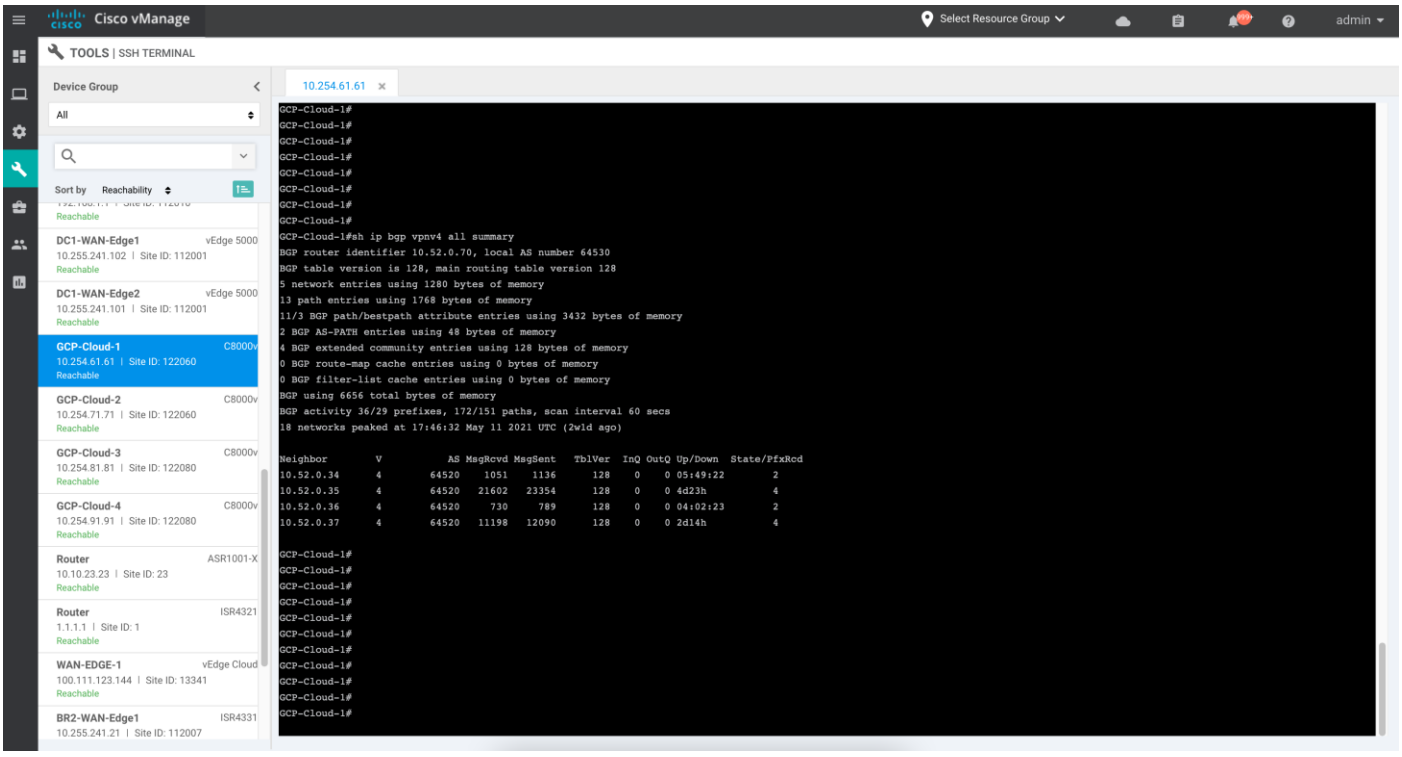

Appendix A: New in this Guide

This guide is new and is not updated from a previous version.

© 2020 Cisco and/or its affiliates. All rights reserved.

Appendix B: Hardware and Software Used for Validation

This guide was validated using the following hardware and software.

Table 3. System Feature Template Settings

Functional Area	Product	Software Version
Cloud	Cisco vManage NMS	20.5
Cloud	Cisco vBond Controller	20.5
Cloud	Cisco vSmart Controller	20.5
Cloud Gateway Devices	Catalyst 8000v	17.5

Appendix C: Catalyst 8000v Configuration

The following section lists out an example Catalyst 8000v configuration.

Cloud Gateway 1 - US-West1

1st Catalyst 8000v in CGW - US-WEST-1

```
system
 system-ip             10.254.61.61
 overlay-id            1
 site-id               122060
 port-offset           1
 control-session-pps   300
 admin-tech-on-failure
 sp-organization-name  "ENB-Solutions - 21615"
 organization-name     "ENB-Solutions - 21615"
 port-hop
 track-transport
 track-default-gateway
 console-baud-rate     19200
 no on-demand enable
 on-demand idle-timeout 10
 vbond vbond204.cisco.com port 12346
!
service tcp-keepalives-in
service tcp-keepalives-out
no service tcp-small-servers
no service udp-small-servers
hostname GCP-Cloud-1
username admin privilege 15 secret 9
$9$3V6L3V6L2VUI2k$ysPnXOdg8RLj9KgMdmfHdSHkdaMmiHzGaUpcqH6pfTo
vrf definition 1
 rd 1:1
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
 address-family ipv4
  route-target export 64530:1
  route-target import 64530:1
  exit-address-family
 !
 address-family ipv6
  exit-address-family
 !
!
vrf definition Mgmt-intf
 rd 1:512
 address-family ipv4
  route-target export 64530:512
  route-target import 64530:512
  exit-address-family
 !
 address-family ipv6
  exit-address-family
 !
!
ip arp proxy disable
no ip finger
no ip rcmd rcp-enable
no ip rcmd rsh-enable
ip as-path access-list 15 permit ^645[2][0-9]$
ip as-path access-list 25 permit .*
no ip dhcp use class
ip host vbond204.cisco.com 52.156.128.118
ip prefix-list GCP_CSR_PREFIX_LIST_POLICY seq 5 permit 10.52.0.70/32
ip bootp server
no ip source-route
no ip http server
no ip http secure-server
ip nat settings central-policy
cdp run
interface GigabitEthernet1
 no shutdown
 arp timeout 1200
 ip address dhcp client-id GigabitEthernet1
 no ip redirects
 ip dhcp client default-router distance 1
 ip mtu    1500
 load-interval 30
 mtu           1500
 negotiation auto
exit
interface GigabitEthernet2
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
 no shutdown
 arp timeout 1200
 vrf forwarding 1
 ip address 10.52.0.38 255.255.255.224
 no ip redirects
 ip mtu    1500
 load-interval 30
 mtu           1500
 negotiation auto
exit
interface GigabitEthernet3
 no shutdown
 arp timeout 1200
 ip address 10.52.0.70 255.255.255.224
 no ip redirects
 ip mtu    1500
 load-interval 30
 mtu           1500
 negotiation auto
exit
interface Tunnel1
 no shutdown
 ip unnumbered GigabitEthernet1
 no ip redirects
 ipv6 unnumbered GigabitEthernet1
 no ipv6 redirects
 tunnel source GigabitEthernet1
 tunnel mode sdwan
exit
interface Tunnel3
 no shutdown
 ip unnumbered GigabitEthernet3
 no ip redirects
 ipv6 unnumbered GigabitEthernet3
 no ipv6 redirects
 tunnel source GigabitEthernet3
 tunnel mode sdwan
exit
route-map GCP_CSR_PREFIX_ROUTE_POLICY permit 10
 match ip address prefix-list GCP_CSR_PREFIX_LIST_POLICY
!
route-map GCP_CSR_PREFIX_ROUTE_POLICY deny 65535
!
route-map GCP_CSR_ROUTE_POLICY deny 1
 match as-path 15
!
route-map GCP_CSR_ROUTE_POLICY permit 11
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
 match as-path 25
!
route-map GCP_CSR_ROUTE_POLICY deny 65535
!
clock timezone UTC 0 0
logging persistent size 104857600 filesize 10485760
logging console
aaa authentication login default local
aaa authorization exec default local
aaa server radius dynamic-author
!
router bgp 64530
 bgp log-neighbor-changes
 neighbor 10.52.0.66 remote-as 64521
 neighbor 10.52.0.66 ebgp-multihop 255
 neighbor 10.52.0.67 remote-as 64521
 neighbor 10.52.0.67 ebgp-multihop 255
 neighbor 10.52.0.68 remote-as 64521
 neighbor 10.52.0.68 ebgp-multihop 255
 neighbor 10.52.0.69 remote-as 64521
 neighbor 10.52.0.69 ebgp-multihop 255
 address-family ipv4 unicast vrf 1
  distance bgp 20 200 20
  neighbor 10.52.0.34 remote-as 64520
  neighbor 10.52.0.34 activate
  neighbor 10.52.0.34 ebgp-multihop 255
  neighbor 10.52.0.34 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.34 send-community both
  neighbor 10.52.0.35 remote-as 64520
  neighbor 10.52.0.35 activate
  neighbor 10.52.0.35 ebgp-multihop 255
  neighbor 10.52.0.35 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.35 send-community both
  neighbor 10.52.0.36 remote-as 64520
  neighbor 10.52.0.36 activate
  neighbor 10.52.0.36 ebgp-multihop 255
  neighbor 10.52.0.36 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.36 send-community both
  neighbor 10.52.0.37 remote-as 64520
  neighbor 10.52.0.37 activate
  neighbor 10.52.0.37 ebgp-multihop 255
  neighbor 10.52.0.37 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.37 send-community both
  propagate-aspath
  redistribute omp
  exit-address-family
 !
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
 address-family ipv4 unicast
  distance bgp 20 200 20
  neighbor 10.52.0.66 activate
  neighbor 10.52.0.66 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.66 send-community both
  neighbor 10.52.0.67 activate
  neighbor 10.52.0.67 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.67 send-community both
  neighbor 10.52.0.68 activate
  neighbor 10.52.0.68 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.68 send-community both
  neighbor 10.52.0.69 activate
  neighbor 10.52.0.69 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.69 send-community both
  redistribute connected
  redistribute connected route-map GCP_CSR_PREFIX_ROUTE_POLICY
  exit-address-family
 !
 timers bgp 60 180
!
snmp-server ifindex persist
line aux 0
 stopbits 1
!
line con 0
 speed    19200
 stopbits 1
!
line vty 0 4
 transport input ssh
!
line vty 5 80
 transport input ssh
!
lldp run
nat64 translation timeout tcp 60
nat64 translation timeout udp 1
sdwan
 interface GigabitEthernet1
  tunnel-interface
   encapsulation ipsec weight 1
   no border
   color biz-internet
   no last-resort-circuit
   no low-bandwidth-link
   no vbond-as-stun-server
   vmanage-connection-preference 5
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
  port-hop
  carrier                       default
  nat-refresh-interval          5
  hello-interval                1000
  hello-tolerance               12
  allow-service all
  no allow-service bgp
  allow-service dhcp
  allow-service dns
  allow-service icmp
  allow-service sshd
  no allow-service netconf
  no allow-service ntp
  no allow-service ospf
  no allow-service stun
  allow-service https
  no allow-service snmp
  no allow-service bfd
 exit
exit
interface GigabitEthernet3
 tunnel-interface
  encapsulation ipsec weight 1
  no border
  color private1
  no last-resort-circuit
  no low-bandwidth-link
  max-control-connections       0
  no vbond-as-stun-server
  vmanage-connection-preference 0
  port-hop
  carrier                       default
  nat-refresh-interval          5
  hello-interval                1000
  hello-tolerance               12
  allow-service all
  allow-service bgp
  no allow-service dhcp
  allow-service dns
  allow-service icmp
  allow-service sshd
  no allow-service netconf
  no allow-service ntp
  no allow-service ospf
  no allow-service stun
  allow-service https
  no allow-service snmp
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
   no allow-service bfd
  exit
 exit
 appqoe
  no tcpopt enable
  no dreopt enable
 !
 omp
  no shutdown
  send-path-limit  4
  ecmp-limit       4
  graceful-restart
  no as-dot-notation
  timers
   holdtime               60
   advertisement-interval 1
   graceful-restart-timer 43200
   eor-timer              300
  exit
  address-family ipv4
   advertise bgp
   advertise connected
   advertise static
  !
  address-family ipv6
   advertise bgp
   advertise connected
   advertise static
  !
 !
!
licensing config enable false
licensing config privacy hostname false
licensing config privacy version false
licensing config utility utility-enable false
bfd color lte
 hello-interval 1000
 no pmtu-discovery
 multiplier     1
!
bfd default-dscp 48
bfd app-route multiplier 2
bfd app-route poll-interval 123400
security
 ipsec
  rekey               86400
  replay-window       512
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
  authentication-type ah-sha1-hmac sha1-hmac
 !
!
sslproxy
 no enable
 rsa-key-modulus      2048
 certificate-lifetime 730
 eckey-type           P256
 ca-tp-label          PROXY-SIGNING-CA
 settings expired-certificate  drop
 settings untrusted-certificate drop
 settings unknown-status       drop
 settings certificate-revocation-check none
 settings unsupported-protocol-versions drop
 settings unsupported-cipher-suites drop
 settings failure-mode         close
 settings minimum-tls-ver      TLSv1
 dual-side optimization enable
!
policy
 no app-visibility
 no app-visibility-ipv6
 no flow-visibility
 no flow-visibility-ipv6
 no implicit-acl-logging
 log-frequency        1000
!
```

2nd Catalyst 8000v – CGW US-West-1

```
system
 system-ip             10.254.71.71
 overlay-id            1
 site-id               122060
 port-offset           1
 control-session-pps   300
 admin-tech-on-failure
 sp-organization-name  "ENB-Solutions - 21615"
 organization-name     "ENB-Solutions - 21615"
 port-hop
 track-transport
 track-default-gateway
 console-baud-rate     19200
 no on-demand enable
 on-demand idle-timeout 10
 vbond vbond204.cisco.com port 12346
!
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
service tcp-keepalives-in
service tcp-keepalives-out
no service tcp-small-servers
no service udp-small-servers
hostname GCP-Cloud-2
username admin privilege 15 secret 9
$9$3V6L3V6L2VUI2k$ysPnXOdg8RLj9KgMdmfHdSHkdaMmiHzGaUpcqH6pfTo
vrf definition 1
 rd 1:1
 address-family ipv4
  route-target export 64530:1
  route-target import 64530:1
  exit-address-family
 !
 address-family ipv6
  exit-address-family
 !
!
vrf definition Mgmt-intf
 rd 1:512
 address-family ipv4
  route-target export 64530:512
  route-target import 64530:512
  exit-address-family
 !
 address-family ipv6
  exit-address-family
 !
!
ip arp proxy disable
no ip finger
no ip rcmd rcp-enable
no ip rcmd rsh-enable
ip as-path access-list 15 permit ^645[2][0-9]$
ip as-path access-list 25 permit .*
no ip dhcp use class
ip host vbond204.cisco.com 52.156.128.118
ip prefix-list GCP_CSR_PREFIX_LIST_POLICY seq 5 permit 10.52.0.198/32
ip bootp server
no ip source-route
no ip http server
no ip http secure-server
ip nat settings central-policy
cdp run
interface GigabitEthernet1
 no shutdown
 arp timeout 1200
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
 ip address dhcp client-id GigabitEthernet1
 no ip redirects
 ip dhcp client default-router distance 1
 ip mtu    1500
 load-interval 30
 mtu           1500
 negotiation auto
exit
interface GigabitEthernet2
 no shutdown
 arp timeout 1200
 vrf forwarding 1
 ip address 10.52.0.166 255.255.255.224
 no ip redirects
 ip mtu    1500
 load-interval 30
 mtu           1500
 negotiation auto
exit
interface GigabitEthernet3
 no shutdown
 arp timeout 1200
 ip address 10.52.0.198 255.255.255.224
 no ip redirects
 ip mtu    1500
 load-interval 30
 mtu           1500
 negotiation auto
exit
interface Tunnel1
 no shutdown
 ip unnumbered GigabitEthernet1
 no ip redirects
 ipv6 unnumbered GigabitEthernet1
 no ipv6 redirects
 tunnel source GigabitEthernet1
 tunnel mode sdwan
exit
interface Tunnel3
 no shutdown
 ip unnumbered GigabitEthernet3
 no ip redirects
 ipv6 unnumbered GigabitEthernet3
 no ipv6 redirects
 tunnel source GigabitEthernet3
 tunnel mode sdwan
exit
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
route-map GCP_CSR_PREFIX_ROUTE_POLICY permit 10
 match ip address prefix-list GCP_CSR_PREFIX_LIST_POLICY
!
route-map GCP_CSR_PREFIX_ROUTE_POLICY deny 65535
!
route-map GCP_CSR_ROUTE_POLICY deny 1
 match as-path 15
!
route-map GCP_CSR_ROUTE_POLICY permit 11
 match as-path 25
!
route-map GCP_CSR_ROUTE_POLICY deny 65535
!
clock timezone UTC 0 0
logging persistent size 104857600 filesize 10485760
logging console
aaa authentication login default local
aaa authorization exec default local
aaa server radius dynamic-author
!
router bgp 64530
 bgp log-neighbor-changes
 neighbor 10.52.0.194 remote-as 64521
 neighbor 10.52.0.194 ebgp-multihop 255
 neighbor 10.52.0.195 remote-as 64521
 neighbor 10.52.0.195 ebgp-multihop 255
 neighbor 10.52.0.196 remote-as 64521
 neighbor 10.52.0.196 ebgp-multihop 255
 neighbor 10.52.0.197 remote-as 64521
 neighbor 10.52.0.197 ebgp-multihop 255
 address-family ipv4 unicast vrf 1
  distance bgp 20 200 20
  neighbor 10.52.0.162 remote-as 64520
  neighbor 10.52.0.162 activate
  neighbor 10.52.0.162 ebgp-multihop 255
  neighbor 10.52.0.162 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.162 send-community both
  neighbor 10.52.0.163 remote-as 64520
  neighbor 10.52.0.163 activate
  neighbor 10.52.0.163 ebgp-multihop 255
  neighbor 10.52.0.163 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.163 send-community both
  neighbor 10.52.0.164 remote-as 64520
  neighbor 10.52.0.164 activate
  neighbor 10.52.0.164 ebgp-multihop 255
  neighbor 10.52.0.164 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.164 send-community both
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
  neighbor 10.52.0.165 remote-as 64520
  neighbor 10.52.0.165 activate
  neighbor 10.52.0.165 ebgp-multihop 255
  neighbor 10.52.0.165 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.165 send-community both
  propagate-aspath
  redistribute omp
  exit-address-family
 !
 address-family ipv4 unicast
  distance bgp 20 200 20
  neighbor 10.52.0.194 activate
  neighbor 10.52.0.194 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.194 send-community both
  neighbor 10.52.0.195 activate
  neighbor 10.52.0.195 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.195 send-community both
  neighbor 10.52.0.196 activate
  neighbor 10.52.0.196 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.196 send-community both
  neighbor 10.52.0.197 activate
  neighbor 10.52.0.197 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.197 send-community both
  redistribute connected
  redistribute connected route-map GCP_CSR_PREFIX_ROUTE_POLICY
  exit-address-family
 !
 timers bgp 60 180
!
snmp-server ifindex persist
line aux 0
 stopbits 1
!
line con 0
 speed    19200
 stopbits 1
!
line vty 0 4
 transport input ssh
!
line vty 5 80
 transport input ssh
!
lldp run
nat64 translation timeout tcp 60
nat64 translation timeout udp 1
sdwan
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
interface GigabitEthernet1
 tunnel-interface
  encapsulation ipsec weight 1
  no border
  color biz-internet
  no last-resort-circuit
  no low-bandwidth-link
  no vbond-as-stun-server
  vmanage-connection-preference 5
  port-hop
  carrier                       default
  nat-refresh-interval          5
  hello-interval                1000
  hello-tolerance               12
  allow-service all
  no allow-service bgp
  allow-service dhcp
  allow-service dns
  allow-service icmp
  allow-service sshd
  no allow-service netconf
  no allow-service ntp
  no allow-service ospf
  no allow-service stun
  allow-service https
  no allow-service snmp
  no allow-service bfd
 exit
exit
interface GigabitEthernet3
 tunnel-interface
  encapsulation ipsec weight 1
  no border
  color private1
  no last-resort-circuit
  no low-bandwidth-link
  max-control-connections       0
  no vbond-as-stun-server
  vmanage-connection-preference 0
  port-hop
  carrier                       default
  nat-refresh-interval          5
  hello-interval                1000
  hello-tolerance               12
  allow-service all
  allow-service bgp
  no allow-service dhcp
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
   allow-service dns
   allow-service icmp
   allow-service sshd
   no allow-service netconf
   no allow-service ntp
   no allow-service ospf
   no allow-service stun
   allow-service https
   no allow-service snmp
   no allow-service bfd
  exit
 exit
 appqoe
  no tcpopt enable
  no dreopt enable
 !
 omp
  no shutdown
  send-path-limit  4
  ecmp-limit       4
  graceful-restart
  no as-dot-notation
  timers
   holdtime               60
   advertisement-interval 1
   graceful-restart-timer 43200
   eor-timer              300
  exit
  address-family ipv4
   advertise bgp
   advertise connected
   advertise static
  !
  address-family ipv6
   advertise bgp
   advertise connected
   advertise static
  !
 !
!
licensing config enable false
licensing config privacy hostname false
licensing config privacy version false
licensing config utility utility-enable false
bfd color lte
 hello-interval 1000
 no pmtu-discovery
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
 multiplier     1
!
bfd default-dscp 48
bfd app-route multiplier 2
bfd app-route poll-interval 123400
security
 ipsec
  rekey               86400
  replay-window       512
  authentication-type ah-sha1-hmac sha1-hmac
 !
!
sslproxy
 no enable
 rsa-key-modulus      2048
 certificate-lifetime 730
 eckey-type           P256
 ca-tp-label          PROXY-SIGNING-CA
 settings expired-certificate  drop
 settings untrusted-certificate drop
 settings unknown-status       drop
 settings certificate-revocation-check none
 settings unsupported-protocol-versions drop
 settings unsupported-cipher-suites drop
 settings failure-mode         close
 settings minimum-tls-ver      TLSv1
 dual-side optimization enable
!
policy
 no app-visibility
 no app-visibility-ipv6
 no flow-visibility
 no flow-visibility-ipv6
 no implicit-acl-logging
 log-frequency        1000
!
```

Cloud Gateway 2 - US-West2

1st Catalyst 8000v - CGW US-West-2

```
system
 system-ip             10.254.81.81
 overlay-id            1
 site-id               122080
 port-offset           1
 control-session-pps   300
 admin-tech-on-failure
 sp-organization-name  "ENB-Solutions - 21615"
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
 organization-name     "ENB-Solutions - 21615"
 port-hop
 track-transport
 track-default-gateway
 console-baud-rate     19200
 no on-demand enable
 on-demand idle-timeout 10
 vbond vbond204.cisco.com port 12346
!
service tcp-keepalives-in
service tcp-keepalives-out
no service tcp-small-servers
no service udp-small-servers
hostname GCP-Cloud-3
username admin privilege 15 secret 9
$9$3V6L3V6L2VUI2k$ysPnXOdg8RLj9KgMdmfHdSHkdaMmiHzGaUpcqH6pfTo
vrf definition 1
 rd 1:1
 address-family ipv4
  route-target export 64530:1
  route-target import 64530:1
  exit-address-family
 !
 address-family ipv6
  exit-address-family
 !
!
vrf definition Mgmt-intf
 rd 1:512
 address-family ipv4
  route-target export 64530:512
  route-target import 64530:512
  exit-address-family
 !
 address-family ipv6
  exit-address-family
 !
!
ip arp proxy disable
no ip finger
no ip rcmd rcp-enable
no ip rcmd rsh-enable
ip as-path access-list 15 permit ^645[2][0-9]$
ip as-path access-list 25 permit .*
no ip dhcp use class
ip host vbond204.cisco.com 52.156.128.118
ip prefix-list GCP_CSR_PREFIX_LIST_POLICY seq 5 permit 10.52.0.71/32
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
ip bootp server
no ip source-route
no ip http server
no ip http secure-server
ip nat settings central-policy
cdp run
interface GigabitEthernet1
 no shutdown
 arp timeout 1200
 ip address dhcp client-id GigabitEthernet1
 no ip redirects
 ip dhcp client default-router distance 1
 ip mtu    1500
 load-interval 30
 mtu           1500
 negotiation auto
exit
interface GigabitEthernet2
 no shutdown
 arp timeout 1200
 vrf forwarding 1
 ip address 10.52.0.39 255.255.255.224
 no ip redirects
 ip mtu    1500
 load-interval 30
 mtu           1500
 negotiation auto
exit
interface GigabitEthernet3
 no shutdown
 arp timeout 1200
 ip address 10.52.0.71 255.255.255.224
 no ip redirects
 ip mtu    1500
 load-interval 30
 mtu           1500
 negotiation auto
exit
interface Tunnel1
 no shutdown
 ip unnumbered GigabitEthernet1
 no ip redirects
 ipv6 unnumbered GigabitEthernet1
 no ipv6 redirects
 tunnel source GigabitEthernet1
 tunnel mode sdwan
exit
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
interface Tunnel3
 no shutdown
 ip unnumbered GigabitEthernet3
 no ip redirects
 ipv6 unnumbered GigabitEthernet3
 no ipv6 redirects
 tunnel source GigabitEthernet3
 tunnel mode sdwan
exit
route-map GCP_CSR_PREFIX_ROUTE_POLICY permit 10
 match ip address prefix-list GCP_CSR_PREFIX_LIST_POLICY
!
route-map GCP_CSR_PREFIX_ROUTE_POLICY deny 65535
!
route-map GCP_CSR_ROUTE_POLICY deny 1
 match as-path 15
!
route-map GCP_CSR_ROUTE_POLICY permit 11
 match as-path 25
!
route-map GCP_CSR_ROUTE_POLICY deny 65535
!
clock timezone UTC 0 0
logging persistent size 104857600 filesize 10485760
logging console
aaa authentication login default local
aaa authorization exec default local
aaa server radius dynamic-author
!
router bgp 64530
 bgp log-neighbor-changes
 neighbor 10.52.0.66 remote-as 64521
 neighbor 10.52.0.66 ebgp-multihop 255
 neighbor 10.52.0.67 remote-as 64521
 neighbor 10.52.0.67 ebgp-multihop 255
 neighbor 10.52.0.68 remote-as 64521
 neighbor 10.52.0.68 ebgp-multihop 255
 neighbor 10.52.0.69 remote-as 64521
 neighbor 10.52.0.69 ebgp-multihop 255
 address-family ipv4 unicast vrf 1
  distance bgp 20 200 20
  neighbor 10.52.0.34 remote-as 64520
  neighbor 10.52.0.34 activate
  neighbor 10.52.0.34 ebgp-multihop 255
  neighbor 10.52.0.34 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.34 send-community both
  neighbor 10.52.0.35 remote-as 64520
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
  neighbor 10.52.0.35 activate
  neighbor 10.52.0.35 ebgp-multihop 255
  neighbor 10.52.0.35 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.35 send-community both
  neighbor 10.52.0.36 remote-as 64520
  neighbor 10.52.0.36 activate
  neighbor 10.52.0.36 ebgp-multihop 255
  neighbor 10.52.0.36 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.36 send-community both
  neighbor 10.52.0.37 remote-as 64520
  neighbor 10.52.0.37 activate
  neighbor 10.52.0.37 ebgp-multihop 255
  neighbor 10.52.0.37 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.37 send-community both
  propagate-aspath
  redistribute omp
  exit-address-family
 !
 address-family ipv4 unicast
  distance bgp 20 200 20
  neighbor 10.52.0.66 activate
  neighbor 10.52.0.66 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.66 send-community both
  neighbor 10.52.0.67 activate
  neighbor 10.52.0.67 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.67 send-community both
  neighbor 10.52.0.68 activate
  neighbor 10.52.0.68 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.68 send-community both
  neighbor 10.52.0.69 activate
  neighbor 10.52.0.69 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.69 send-community both
  redistribute connected
  redistribute connected route-map GCP_CSR_PREFIX_ROUTE_POLICY
  exit-address-family
 !
 timers bgp 60 180
!
snmp-server ifindex persist
line aux 0
 stopbits 1
!
line con 0
 speed    19200
 stopbits 1
!
line vty 0 4
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
 transport input ssh
!
line vty 5 80
 transport input ssh
!
lldp run
nat64 translation timeout tcp 60
nat64 translation timeout udp 1
sdwan
 interface GigabitEthernet1
  tunnel-interface
   encapsulation ipsec weight 1
   no border
   color biz-internet
   no last-resort-circuit
   no low-bandwidth-link
   no vbond-as-stun-server
   vmanage-connection-preference 5
   port-hop
   carrier                       default
   nat-refresh-interval          5
   hello-interval                1000
   hello-tolerance               12
   allow-service all
   no allow-service bgp
   allow-service dhcp
   allow-service dns
   allow-service icmp
   allow-service sshd
   no allow-service netconf
   no allow-service ntp
   no allow-service ospf
   no allow-service stun
   allow-service https
   no allow-service snmp
   no allow-service bfd
  exit
 exit
 interface GigabitEthernet3
  tunnel-interface
   encapsulation ipsec weight 1
   no border
   color private1
   no last-resort-circuit
   no low-bandwidth-link
   max-control-connections       0
   no vbond-as-stun-server
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
  vmanage-connection-preference 0
  port-hop
  carrier                       default
  nat-refresh-interval          5
  hello-interval                1000
  hello-tolerance               12
  allow-service all
  allow-service bgp
  no allow-service dhcp
  allow-service dns
  allow-service icmp
  allow-service sshd
  no allow-service netconf
  no allow-service ntp
  no allow-service ospf
  no allow-service stun
  allow-service https
  no allow-service snmp
  no allow-service bfd
 exit
exit
appqoe
 no tcpopt enable
 no dreopt enable
!
omp
 no shutdown
 send-path-limit  4
 ecmp-limit       4
 graceful-restart
 no as-dot-notation
 timers
  holdtime               60
  advertisement-interval 1
  graceful-restart-timer 43200
  eor-timer              300
 exit
 address-family ipv4
  advertise bgp
  advertise connected
  advertise static
 !
 address-family ipv6
  advertise bgp
  advertise connected
  advertise static
 !
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
 !
!
licensing config enable false
licensing config privacy hostname false
licensing config privacy version false
licensing config utility utility-enable false
bfd color lte
 hello-interval 1000
 no pmtu-discovery
 multiplier     1
!
bfd default-dscp 48
bfd app-route multiplier 2
bfd app-route poll-interval 123400
security
 ipsec
  rekey               86400
  replay-window       512
  authentication-type ah-sha1-hmac sha1-hmac
 !
!
sslproxy
 no enable
 rsa-key-modulus      2048
 certificate-lifetime 730
 eckey-type           P256
 ca-tp-label          PROXY-SIGNING-CA
 settings expired-certificate  drop
 settings untrusted-certificate drop
 settings unknown-status       drop
 settings certificate-revocation-check none
 settings unsupported-protocol-versions drop
 settings unsupported-cipher-suites drop
 settings failure-mode         close
 settings minimum-tls-ver      TLSv1
 dual-side optimization enable
!
policy
 no app-visibility
 no app-visibility-ipv6
 no flow-visibility
 no flow-visibility-ipv6
 no implicit-acl-logging
 log-frequency        1000
!
```

© 2020 Cisco and/or its affiliates. All rights reserved.

2st Catalyst 8000v – CGW US-West-2

```
system
 system-ip             10.254.91.91
 overlay-id            1
 site-id               122080
 port-offset           1
 control-session-pps   300
 admin-tech-on-failure
 sp-organization-name  "ENB-Solutions - 21615"
 organization-name     "ENB-Solutions - 21615"
 port-hop
 track-transport
 track-default-gateway
 console-baud-rate     19200
 no on-demand enable
 on-demand idle-timeout 10
 vbond vbond204.cisco.com port 12346
!
service tcp-keepalives-in
service tcp-keepalives-out
no service tcp-small-servers
no service udp-small-servers
hostname GCP-Cloud-4
username admin privilege 15 secret 9
$9$3V6L3V6L2VUI2k$ysPnXOdg8RLj9KgMdmfHdSHkdaMmiHzGaUpcqH6pfTo
vrf definition 1
 rd 1:1
 address-family ipv4
  route-target export 64530:1
  route-target import 64530:1
  exit-address-family
 !
 address-family ipv6
  exit-address-family
 !
!
vrf definition Mgmt-intf
 rd 1:512
 address-family ipv4
  route-target export 64530:512
  route-target import 64530:512
  exit-address-family
 !
 address-family ipv6
  exit-address-family
 !
!
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
ip arp proxy disable
no ip finger
no ip rcmd rcp-enable
no ip rcmd rsh-enable
ip as-path access-list 15 permit ^645[2][0-9]$
ip as-path access-list 25 permit .*
no ip dhcp use class
ip host vbond204.cisco.com 52.156.128.118
ip prefix-list GCP_CSR_PREFIX_LIST_POLICY seq 5 permit 10.52.0.199/32
ip bootp server
no ip source-route
no ip http server
no ip http secure-server
ip nat settings central-policy
cdp run
interface GigabitEthernet1
 no shutdown
 arp timeout 1200
 ip address dhcp client-id GigabitEthernet1
 no ip redirects
 ip dhcp client default-router distance 1
 ip mtu    1500
 load-interval 30
 mtu           1500
 negotiation auto
exit
interface GigabitEthernet2
 no shutdown
 arp timeout 1200
 vrf forwarding 1
 ip address 10.52.0.167 255.255.255.224
 no ip redirects
 ip mtu    1500
 load-interval 30
 mtu           1500
 negotiation auto
exit
interface GigabitEthernet3
 no shutdown
 arp timeout 1200
 ip address 10.52.0.199 255.255.255.224
 no ip redirects
 ip mtu    1500
 load-interval 30
 mtu           1500
 negotiation auto
exit
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
interface Tunnel1
 no shutdown
 ip unnumbered GigabitEthernet1
 no ip redirects
 ipv6 unnumbered GigabitEthernet1
 no ipv6 redirects
 tunnel source GigabitEthernet1
 tunnel mode sdwan
exit
interface Tunnel3
 no shutdown
 ip unnumbered GigabitEthernet3
 no ip redirects
 ipv6 unnumbered GigabitEthernet3
 no ipv6 redirects
 tunnel source GigabitEthernet3
 tunnel mode sdwan
exit
route-map GCP_CSR_PREFIX_ROUTE_POLICY permit 10
 match ip address prefix-list GCP_CSR_PREFIX_LIST_POLICY
!
route-map GCP_CSR_PREFIX_ROUTE_POLICY deny 65535
!
route-map GCP_CSR_ROUTE_POLICY deny 1
 match as-path 15
!
route-map GCP_CSR_ROUTE_POLICY permit 11
 match as-path 25
!
route-map GCP_CSR_ROUTE_POLICY deny 65535
!
clock timezone UTC 0 0
logging persistent size 104857600 filesize 10485760
logging console
aaa authentication login default local
aaa authorization exec default local
aaa server radius dynamic-author
!
router bgp 64530
 bgp log-neighbor-changes
 neighbor 10.52.0.194 remote-as 64521
 neighbor 10.52.0.194 ebgp-multihop 255
 neighbor 10.52.0.195 remote-as 64521
 neighbor 10.52.0.195 ebgp-multihop 255
 neighbor 10.52.0.196 remote-as 64521
 neighbor 10.52.0.196 ebgp-multihop 255
 neighbor 10.52.0.197 remote-as 64521
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
 neighbor 10.52.0.197 ebgp-multihop 255
 address-family ipv4 unicast vrf 1
  distance bgp 20 200 20
  neighbor 10.52.0.162 remote-as 64520
  neighbor 10.52.0.162 activate
  neighbor 10.52.0.162 ebgp-multihop 255
  neighbor 10.52.0.162 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.162 send-community both
  neighbor 10.52.0.163 remote-as 64520
  neighbor 10.52.0.163 activate
  neighbor 10.52.0.163 ebgp-multihop 255
  neighbor 10.52.0.163 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.163 send-community both
  neighbor 10.52.0.164 remote-as 64520
  neighbor 10.52.0.164 activate
  neighbor 10.52.0.164 ebgp-multihop 255
  neighbor 10.52.0.164 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.164 send-community both
  neighbor 10.52.0.165 remote-as 64520
  neighbor 10.52.0.165 activate
  neighbor 10.52.0.165 ebgp-multihop 255
  neighbor 10.52.0.165 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.165 send-community both
  propagate-aspath
  redistribute omp
  exit-address-family
 !
 address-family ipv4 unicast
  distance bgp 20 200 20
  neighbor 10.52.0.194 activate
  neighbor 10.52.0.194 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.194 send-community both
  neighbor 10.52.0.195 activate
  neighbor 10.52.0.195 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.195 send-community both
  neighbor 10.52.0.196 activate
  neighbor 10.52.0.196 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.196 send-community both
  neighbor 10.52.0.197 activate
  neighbor 10.52.0.197 route-map GCP_CSR_ROUTE_POLICY out
  neighbor 10.52.0.197 send-community both
  redistribute connected
  redistribute connected route-map GCP_CSR_PREFIX_ROUTE_POLICY
  exit-address-family
 !
 timers bgp 60 180
!
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
snmp-server ifindex persist
line aux 0
 stopbits 1
!
line con 0
 speed    19200
 stopbits 1
!
line vty 0 4
 transport input ssh
!
line vty 5 80
 transport input ssh
!
lldp run
nat64 translation timeout tcp 60
nat64 translation timeout udp 1
sdwan
 interface GigabitEthernet1
  tunnel-interface
   encapsulation ipsec weight 1
   no border
   color biz-internet
   no last-resort-circuit
   no low-bandwidth-link
   no vbond-as-stun-server
   vmanage-connection-preference 5
   port-hop
   carrier                       default
   nat-refresh-interval          5
   hello-interval                1000
   hello-tolerance               12
   allow-service all
   no allow-service bgp
   allow-service dhcp
   allow-service dns
   allow-service icmp
   allow-service sshd
   no allow-service netconf
   no allow-service ntp
   no allow-service ospf
   no allow-service stun
   allow-service https
   no allow-service snmp
   no allow-service bfd
  exit
 exit
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
interface GigabitEthernet3
 tunnel-interface
  encapsulation ipsec weight 1
  no border
  color private1
  no last-resort-circuit
  no low-bandwidth-link
  max-control-connections       0
  no vbond-as-stun-server
  vmanage-connection-preference 0
  port-hop
  carrier                       default
  nat-refresh-interval          5
  hello-interval                1000
  hello-tolerance               12
  allow-service all
  allow-service bgp
  no allow-service dhcp
  allow-service dns
  allow-service icmp
  allow-service sshd
  no allow-service netconf
  no allow-service ntp
  no allow-service ospf
  no allow-service stun
  allow-service https
  no allow-service snmp
  no allow-service bfd
 exit
exit
appqoe
 no tcpopt enable
 no dreopt enable
!
omp
 no shutdown
 send-path-limit  4
 ecmp-limit       4
 graceful-restart
 no as-dot-notation
 timers
  holdtime               60
  advertisement-interval 1
  graceful-restart-timer 43200
  eor-timer              300
 exit
 address-family ipv4
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
   advertise bgp
   advertise connected
   advertise static
  !
  address-family ipv6
   advertise bgp
   advertise connected
   advertise static
  !
 !
!
licensing config enable false
licensing config privacy hostname false
licensing config privacy version false
licensing config utility utility-enable false
bfd color lte
 hello-interval 1000
 no pmtu-discovery
 multiplier     1
!
bfd default-dscp 48
bfd app-route multiplier 2
bfd app-route poll-interval 123400
security
 ipsec
  rekey               86400
  replay-window       512
  authentication-type ah-sha1-hmac sha1-hmac
 !
!
sslproxy
 no enable
 rsa-key-modulus      2048
 certificate-lifetime 730
 eckey-type           P256
 ca-tp-label          PROXY-SIGNING-CA
 settings expired-certificate  drop
 settings untrusted-certificate drop
 settings unknown-status       drop
 settings certificate-revocation-check none
 settings unsupported-protocol-versions drop
 settings unsupported-cipher-suites drop
 settings failure-mode         close
 settings minimum-tls-ver      TLSv1
 dual-side optimization enable
!
policy
```

© 2020 Cisco and/or its affiliates. All rights reserved.

```
 no app-visibility
 no app-visibility-ipv6
 no flow-visibility
 no flow-visibility-ipv6
 no implicit-acl-logging
 log-frequency        1000
!
```

Appendix D: Glossary

VPN Virtual Private Network

CGW Cloud Gateway

VPC Virtual Private Cloud

WAN Wide Area Network

DNS Domain Name Server

GW Gateway

VM Virtual Machine

Feedback

For comments and suggestions about this guide and related guides, join the discussion on Cisco Community at https://cs.co/en-cvds.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

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: https://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. (1110R)

© 2020 Cisco and/or its affiliates. All rights reserved.