

Aviatrix Secure Hybrid Connectivity with GCP Network Connectivity Center

Private datacenter access to Google APIs & private service-connect endpoints

Reference Architecture

| | |
|--|----|
| Preface..... | 1 |
| Guide Types..... | 2 |
| Disclaimer..... | 2 |
| Getting the Latest Versions..... | 2 |
| Purpose of This Guide..... | 2 |
| Audience..... | 3 |
| Introduction..... | 3 |
| Google Network Connectivity Center Overview..... | 4 |
| Hybrid spokes..... | 5 |
| Router appliance spokes..... | 6 |
| Network Architecture Overview..... | 6 |
| Design Principles..... | 7 |
| Design Goals..... | 7 |
| Core Design Concepts..... | 8 |
| Network Connectivity Center..... | 8 |
| Google APIs..... | 8 |
| Aviatrix Transit Firenet Gateways..... | 12 |
| High Availability and Resilience..... | 12 |
| Aviatrix Edge..... | 13 |
| Scalability and Performance..... | 13 |
| Deployment Strategies..... | 14 |
| Configuration Steps..... | 14 |
| Deploy Aviatrix Transit Gateways..... | 14 |
| Associate Next-Generation Firewalls with Aviatrix Firenet..... | 15 |

| | |
|--|----|
| Deploy GCP Network Connectivity Hub (NCC). | 16 |
| Deploy GCP NCC VPC Spoke | 17 |
| Deploy GCP NCC Router Appliance Spokes | 18 |
| Deploy GCP Cloud Routers..... | 20 |
| Configure VPC Static Routes..... | 23 |
| Configure BGP Peerings | 24 |
| Configure Aviatrix Firenet Inspection | 24 |
| Deploy Google Kubernetes Engine | 25 |
| Deploy an Aviatrix Gateway for SNAT/DNAT. | 25 |
| Configure GCP Firewall rules for Health Checks. | 27 |
| Deploy a GCP Load Balancer | 27 |
| Establish DNS | 30 |
| Using round-robin DNS vs. GCP Load Balancer (optional) | 31 |
| Design Validation | 32 |
| Monitoring and Management Best Practices..... | 32 |

Preface

Guide Types



Overview guides offer high-level introductions to technologies or concepts with an emphasis on business value.

Technology guides provide introduction to product capabilities for using Aviatrix to provide visibility, control, and protection to applications built in a specific environment. These guides describe the technologies providing examples and should be considered required reading prior to using their companion design & deployment guides.

Design & Deployment guides provide definitive guidance for different deployment scenarios, as well as procedures for combining Aviatrix technologies with third-party technologies in an integrated design. To ensure the design is reproducible, these guides provide Terraform templates for deployment of Aviatrix & third-party vendors such as CSPs. Will include best practices recommendations.

Disclaimer

The guides occasionally describe products from other companies. While the steps and screenshots were accurate at the time of publication, those companies may have since updated their user interfaces, processes, or requirements. Please refer to the external vendor website for the latest documentation.

Change History

| Document Version | Aviatrix Controller Versions | Date Changed | Modified By | Change Log |
|------------------|------------------------------|--------------|-----------------------|-------------------------|
| 1.0 | 7.2 and 8.0 | June 2025 | Solution Architecture | First published version |

Getting the Latest Versions

Access the latest reference architecture guides at: <https://architecture.aviatrix.com>

Purpose of This Guide

This guide explains how to use Aviatrix to enhance visibility and securely connect applications on Google Cloud to the enterprise data center.

This guide:

- Provides an overview of how Aviatrix helps organizations manage security risks and compliance when connecting their data center to public cloud infrastructure.
- Links the technical design aspects of Google Cloud and Aviatrix, exploring design variations.

This guide offers a set of decision criteria for various deployment scenarios, along with detailed procedures for enabling features within Google Cloud and Aviatrix to achieve a cohesive and integrated design.

Audience

This guide is for technical readers, including system architects and design engineers, who want to leverage Aviatrix to securely connect Google cloud and the enterprise data center. This guide assumes the reader is familiar with the basic concepts of applications, networking, security, and high availability (HA). The reader should also possess a basic understanding of network and data center architectures. To be successful, you must have a working knowledge of the Aviatrix platform.

Introduction

This guide provides a reference design for securely connecting the enterprise data center to Google Cloud (GCP) using the Google Network Connectivity Center (NCC) and Aviatrix. It also offers design options that allow private datacenter access to Google APIs and private service-connect endpoints, such as private access to a GKE cluster.

Aviatrix is utilized to deliver high-throughput Layer 3 encryption over the Google Cloud Interconnect used to link the data center and GCP.

The following sections describe the Google Network Connectivity Center and include additional Google documentation URLs for reference. A detailed network design diagram is followed by design principles and decision criteria.

Google Network Connectivity Center Overview

Google Network Connectivity Center is an orchestration framework that simplifies network connectivity among *spoke* resources that are connected to a central management resource called a *hub*. Network Connectivity Center supports three types of spokes:

- Virtual Private Cloud (VPC) spokes
- Producer VPC spokes
- Hybrid spokes, consisting of:
 - HA VPN tunnels
 - Cloud Interconnect VLAN attachments
 - Router appliance VMs

With the hub and spoke connectivity, you can do the following:

- Connect multiple VPC networks to one another. The VPC networks can be located across different projects in the same Google Cloud organization or different organizations.
- Connect multiple VPC networks to on-premise or other cloud provider networks. These external networks can be reached through any type of hybrid spoke. This approach is known as *site-to-cloud connectivity*.
- Use Router appliance VMs to manage connectivity between your VPC networks.
- Use a Google Cloud VPC network as an enterprise-wide area network (WAN) to connect networks that are outside of Google Cloud. You can establish connectivity between your external sites by using any type of hybrid spoke. This approach is known as *site-to-site connectivity*.

Hybrid spokes

A hybrid spoke represents one or more network connectivity resources that are connected to a hub. A hybrid spoke type can be any of the following resources that a spoke is associated with:

- Router appliance VMs
- HA VPN tunnels
- Cloud Interconnect VLAN attachments

A single hybrid spoke can be associated with more than one resource of the same type. For example, a hybrid spoke can reference two or more HA VPN tunnels, but that same hybrid spoke cannot also reference Router appliance VMs or Cloud Interconnect VLAN attachments. A hybrid spoke must be in the same project as the Network Connectivity Center hub.

Site-to-site data transfer using hybrid spokes requires that the spokes be located in the same VPC network.

Router appliance spokes

A spoke associated with a Router appliance VM instance supports the following use cases:

- **IPv4 site-to-cloud connectivity:** Establish connectivity between an external site and your VPC network resources.
- **IPv4 site-to-site data transfer:** Use Google's network as part of a wide area network (WAN) that includes your external sites to move data between all the sites.
- **IPv4 connectivity between VPC networks:** Use a third-party network virtual appliance to establish connectivity between your VPC networks.

All site-to-site spokes that are connected to the same hub must have all of their backing resources in the same VPC network.

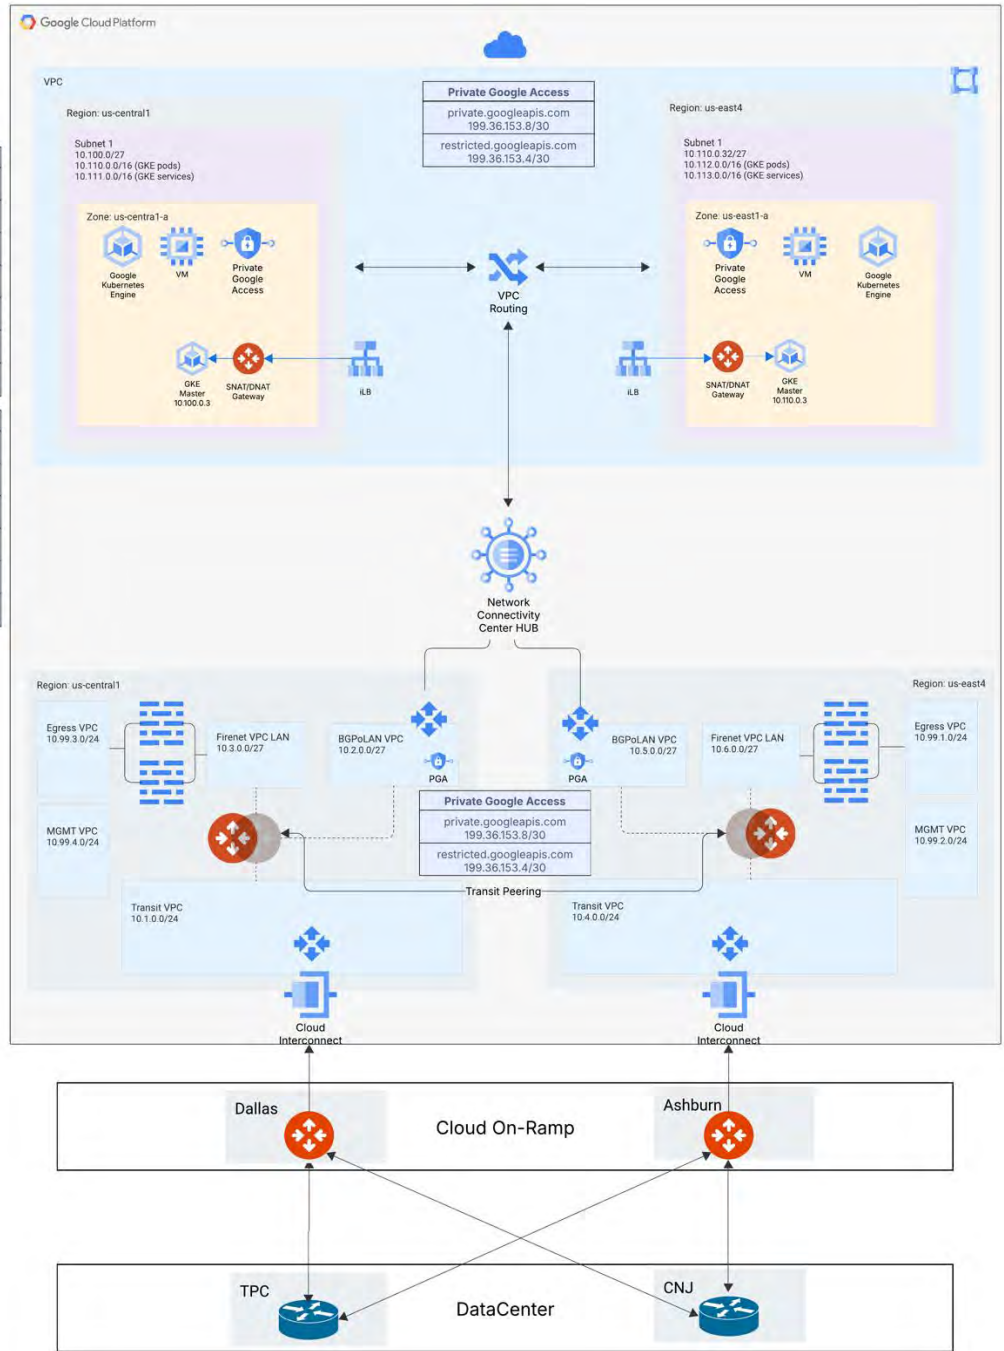
<https://cloud.google.com/vpc/docs/private-service-connect>

<https://cloud.google.com/kubernetes-engine/docs/concepts/private-service-connect>

Network Architecture Overview

| DNAT Policy | |
|------------------|-----------------------|
| Source CIDR | On-premise addresses |
| Destination CIDR | Load Balancer Address |
| Protocol | TCP |
| Interface | Load Balancer Address |
| Connection | On-premise addresses |
| DNAT IPs | GKE Endpoint |
| DNAT Port | 443 |

| SNAT Policy | |
|------------------|-----------------------|
| Source CIDR | On-premise addresses |
| Destination CIDR | GKE Endpoint |
| Protocol | TCP |
| Interface | Load Balancer Address |
| Connection | On-premise addresses |
| SNAT IPs | Aviatix Gateway IP |



Design Principles

Design Goals

The following goals are addressed:

- Connectivity between on-premises and VPCs attached to the NCC hub.
- Traffic from NCC attached VPCs egresses to the Internet via the Aviatrix Transit Firenet security stack with customer-managed next-generation firewalls.
- Simplified egress to the Internet without the need for tags.
- Aviatrix Transit GWs in GCP and Aviatrix Edge GWs that can be deployed in the data center or interconnect providers like Equinix, Megaport, etc. provide the high throughput encryption over Google Interconnect.
- GKE Management endpoint access from on-premises.
- Private access to Google APIs and services from on-premises.
- Overall zonal and regional redundancy to ensure high availability.

Core Design Concepts

Network Connectivity Center

GCP Network Connectivity Center (NCC) is used to connect customer VPCs to the Aviatrix Overlay and dynamically share routing information between the two. An **NCC VPC spoke** is deployed in the host project, with one **router appliance spoke** deployed in the BGPoLAN VPC for each respective region. Three spokes will be deployed in total.

Google APIs

Google's API services can be accessed through two /30 networks from the cloud environment: ***restricted.googleapis.com*** (199.36.153.4/30) and ***private.googleapis.com*** (199.36.153.8/30). These endpoints are accessible by DNS hostname.

Enable **Private Google Access (PGA)** on the BGPoLAN subnets in each region to access GCP Private Google Access endpoints from on-premises.

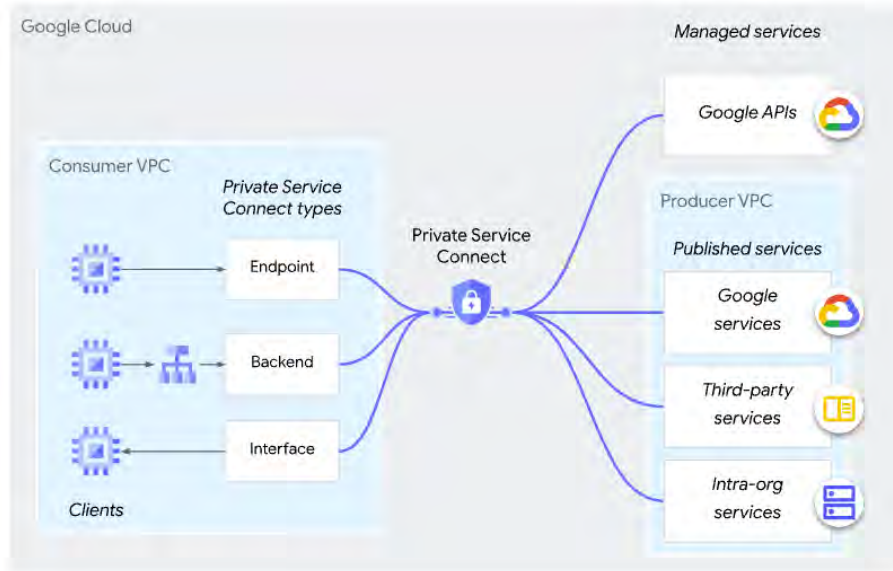
Cloud Routers should be deployed *with custom routes enabled*. In addition to the **"Advertise all subnets visible to the Cloud Router"** setting, custom ranges for the two previously mentioned googleapis networks must be added. Note that [Custom route advertisement in Cloud Router](#) allows manual control over the prefixes that are advertised by hybrid spokes. You can specify custom advertised routes (including default routes, such as 0.0.0.0/0 for IPv4) for all BGP sessions when you don't need automatic advertisement of VPC spoke subnets. By default, other VPC spoke subnets are NOT advertised, which means that on-premises locations don't automatically learn about the reachability to these IP address ranges.

To ensure that traffic to Google APIs favors a single region, *BGP path selection manipulation (such as local preference or AS path prepending) should be implemented on-premises*.

On-premises traffic reaches **Google API endpoints** at Transits without entering the NCC hub.

Subnets in the host project should enable Private Google Access. Traffic from the GCP host project can reach Google API endpoints deployed within the same VPC, without traversing the Aviatrix Transits.

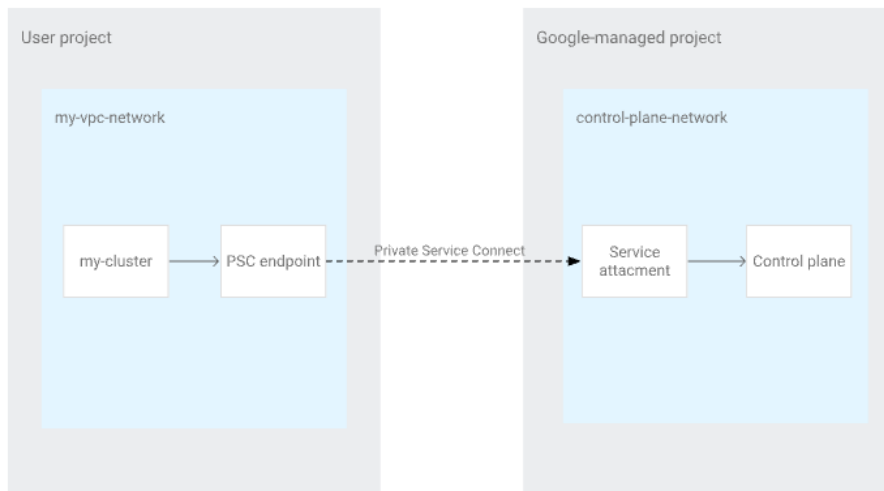
Access model for GCP Private Service Connects:



Google Kubernetes Engine

GCP PaaS services like **Google Kubernetes Engine (GKE)** are hosted in the shared VPC host project. The Kubernetes API server is accessible via an address in the host project subnet. This subnet *IS* advertised through an NCC VPC spoke to the hub, Aviatrix Transits, and on-premises users. However, the endpoint itself is only reachable inside GCP VPCs themselves. Management access CIDRs must be configured to enable required access.

Access model for GKE API server:



GKE endpoints are regional and necessitate the use of an SNAT/DNAT or proxy solution for connectivity. This can be accomplished by using a combination of either a load balancer or round-robin DNS and a non-attached Aviatrix spoke gateway (or similar customer managed NAT or proxy solution) for on-premises access. Direct private access to the endpoint from outside of the VPC (on-premises or other spoke VPCs) is not supported due to internal GCP routing.

GCP load balancers cannot directly utilize this endpoint as a backend. Consequently, the integration of a load balancer with Aviatrix SNAT/DNAT gateway or customer managed NAT or proxy is essential.

Note: Aviatrix preserves the private IP address of gateways when performing image upgrades. The use of round-robin DNS entries (pointing to the Aviatrix spoke gateway pair)

simplifies maintenance of the solution. **If using GCP load balancers, please ensure backend configurations are updated after performing image upgrades.**

Aviatrix Transit Firenet Gateways

Aviatrix Transit Firenet Gateways must be configured to advertise the on-premises routes to the NCC hub when using the Aviatrix NAT solution. **A default route advertisement alone will not suffice due to routing priority and the use of Aviatrix tags for gateway operation.**

Example: The tagged default route for the Aviatrix instance takes priority over the BGP learned default route. Ensure more specific on-premises CIDRs (192.168.0.0/16 in this instance) are advertised to allow traffic to return to end-users.

| Status | Name | Type | IP version | Destination IP range | Priority | Scope limits | Next hop |
|--------|--------------------------------------|-------------|------------|----------------------|----------|---|-----------------------------|
| ✓ | avx-4940b2b09ce847fd8ba780d1f090f0dc | Static | IPv4 | 0.0.0.0/0 | 1000 | Instance tags: avx-gcp-host-project-gbl | Default internet gateway |
| ✓ | ncc-dynamic-route-0862852956 | NCC dynamic | IPv4 | 0.0.0.0/0 | 0 | — | Hub ncc_hub |
| ✓ | ncc-dynamic-route-2746832898 | NCC dynamic | IPv4 | 0.0.0.0/0 | 0 | — | Hub ncc_hub |
| ✓ | ncc-dynamic-route-0994765035 | NCC dynamic | IPv4 | 192.168.0.0/16 | 0 | — | Hub ncc_hub |
| ✓ | ncc-dynamic-route-2376355360 | NCC dynamic | IPv4 | 192.168.0.0/16 | 0 | — | Hub ncc_hub |

High Availability and Resilience

High Availability is achieved by deploying Aviatrix Transit Firenet Gateway resources in at least two cloud regions with primary and secondary members in different zones. Next-generation firewalls should be similarly deployed. If an Aviatrix Transit Gateway fails, the control plane will redirect traffic to a redundant member. Likewise, if a firewall fails, the control plane will redirect traffic to another healthy firewall.

Deploy BGP peerings between Aviatrix Transit Gateways and GCP cloud routers in HA mode, with both primary and HA gateways connecting to regional cloud routers. Ensure GCP cloud routers have redundant interfaces, with each Aviatrix Transit peering using a separate router interface. **Each Aviatrix BGP peering must occur in a unique BGPoLAN subnet for that peering.**

Under normal circumstances, GCP Cloud Routers only advertise routes from their own region. To ensure reachability, configure Aviatrix Transit Gateways in each region. Transit gateways should be peered between regions.

In the event both Aviatrix Transit Gateways in a region fail, GCP Cloud Routers will begin advertising routes from the failed region to Aviatrix in the alternate region. This is referred to as **“automatic cross-regional failover”**.

The GCP Network Connectivity Hub will establish a connection to a Global Shared VPC via an NCC VPC-based spoke.

Aviatrix Edge

Aviatrix Edge connectivity from interconnect providers must be redundant with a minimum of two connections from each colocation center. Aviatrix Edge should be deployed in pairs for redundancy.

Aviatrix Edge can be deployed as either [Spoke](#) or [Transit](#). If Aviatrix overlay connectivity between data centers is necessary, deploy Edge as Transit (EaT) and configure Transit to Transit peering. If this connectivity is not required, or handled by other means (dedicated interconnects, etc.), deploy Aviatrix Edge as a Spoke. (EaS). Deployment may use either physical Aviatrix Edge Appliances (AEP), Edge on customer-managed virtualization, or Edge from marketplace providers like Equinix or Megaport.

Advanced features like Aviatrix Distributed Cloud Firewall are only available on Spokes. For pricing differences between EaS and EaT deployment, consult your account team.

Edge is designed to provide failover abilities by means of BGP or VRRP.

On-premises access to the cloud should follow standard network connectivity practices, including redundant fault-tolerant circuits, diverse points of entry, and the use of multiple carriers.

Scalability and Performance

Sizing Aviatrix Transits depends on your environment's traffic needs. [Use the Aviatrix Gateway Sizing Best Practices Guide for capacity recommendations.](#) For GCP, choose instance sizes comparable to those of other providers. QA to confirm: Routes, BGPoLAN advertisements.

Cloud routers are limited to learning 5,000 BGP prefixes. The BGP session will be reset if it receives more than this limit. Each Cloud router is limited to 200 custom advertised routes per BGP session. For a comprehensive list of cloud router quotas and limitations see the [official GCP documentation.](#)

Deployment Strategies

A Terraform module for this deployment will be released shortly and available in the [Aviatrix Github Module Repository.](#)

Configuration Steps

Deploy Aviatrix Transit Gateways

- Deploy a pair in each region with Firenet and BGPoLAN enabled.
- Aviatrix Transit Firenet requires the following VPCs each with at least 2 subnets in different zones:
 - LAN
 - Transit Firenet
 - Egress
 - Management
 - BGPoLAN
- Manually advertise 0.0.0.0/0 and on-premises CIDRs via Aviatrix BGP from the Transit Gateways to each Cloud Router (on a connection basis, not globally)

Manual BGP Advertisements:

The screenshot shows the AWS Transit Gateway console for a resource named 'avx-transit-central'. The navigation tabs include 'Gateways', 'Overview', 'Transit Gateways', 'Spoke Gateways', and 'Specialty Gateways'. The 'Transit Gateways' tab is active, and the 'Details' sub-tab is selected. The configuration is organized into sections: 'General', 'Border Gateway Protocol (BGP)', and 'Manual Advertised CIDRs'. The 'BGP' section includes a toggle for 'Preserve AS Path' (currently Off) and a text input for 'Local ASN' with the value '65301'. The 'Manual Advertised CIDRs' section has two sub-sections: 'Advertised CIDRs (Per Gateway)' which is currently empty, and 'External Connection' which has a dropdown menu showing 'avx-transit-central-cr-ncc-to-avx'. Below this, the 'Advertised CIDRs (Per Connection)' field contains the value '0.0.0.0,192.168.0.0/16'.

Gateways Overview **Transit Gateways** Spoke Gateways Specialty Gateways G

< avx-transit-central

Details Instances Attachments VPC/VNet Route Tables Gateway Routes Int

General

Border Gateway Protocol (BGP)

Preserve AS Path Off

Local ASN

65301

Manual Advertised CIDRs

Advertised CIDRs (Per Gateway)

External Connection

avx-transit-central-cr-ncc-to-avx

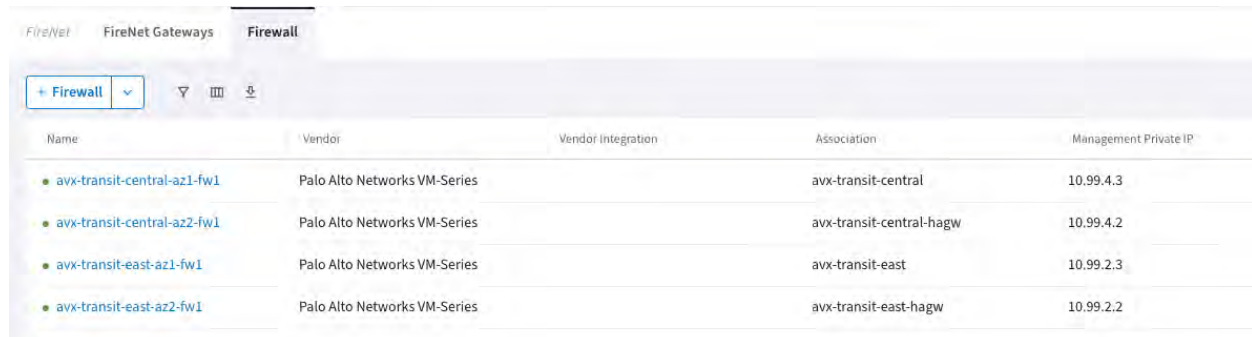
Advertised CIDRs (Per Connection)

0.0.0.0,192.168.0.0/16

Associate Next-Generation Firewalls with Aviatrix Firenet.

- Associate a pair of Next-Generation Firewalls with Aviatrix Firenets in each region, following the [Aviatrix Transit Firenet Workflow](#).
- The configuration of policy is outside of the scope of this document.

Aviatrix Firenet/Firewall Association:



| Name | Vendor | Vendor Integration | Association | Management Private IP |
|-----------------------------|------------------------------|--------------------|--------------------------|-----------------------|
| avx-transit-central-az1-fw1 | Palo Alto Networks VM-Series | | avx-transit-central | 10.99.4.3 |
| avx-transit-central-az2-fw1 | Palo Alto Networks VM-Series | | avx-transit-central-hagw | 10.99.4.2 |
| avx-transit-east-az1-fw1 | Palo Alto Networks VM-Series | | avx-transit-east | 10.99.2.3 |
| avx-transit-east-az2-fw1 | Palo Alto Networks VM-Series | | avx-transit-east-hagw | 10.99.2.2 |

Deploy GCP Network Connectivity Hub (NCC).

- Deploy the GCP Network connectivity hub.
- Enable “Export Private Service Connects”.

Network Connectivity Center Hub Creation:

The screenshot shows the 'Hub details' page for a hub named 'ncc_hub'. The breadcrumb navigation is 'Network Connectivity / Hubs / Hub: ncc_hub'. The page title is 'Hub details' with a 'DELETE HUB' button. The hub name 'ncc_hub' is displayed, with tabs for 'HUB', 'SPOKES', and 'ROUTES'. An 'EDIT HUB' button is visible. The 'Description' field contains the text 'Created by terraform-aviatrix-gcp-ncc module.'. The 'Policy mode' is set to 'Preset topology' and the 'Preset topology' is 'Mesh topology'. A key-value pair is shown: 'Key 1' is 'goog-terraform-provisioned' and 'Value 1' is 'true'. There is an '+ ADD LABEL' button. The 'Private Service Connect connection propagation' section is expanded, showing a description and a 'Learn more' link. The 'On' radio button is selected. At the bottom, there are 'SAVE' and 'CANCEL' buttons.

Network Connectivity / Hubs / Hub: ncc_hub

Hub details DELETE HUB

ncc_hub

HUB SPOKES ROUTES

EDIT HUB

Description
Created by terraform-aviatrix-gcp-ncc module.

Policy mode
Preset topology

Preset topology
Mesh topology

Key 1
goog-terraform-provisioned

Value 1
true

+ ADD LABEL

Private Service Connect connection propagation

Private Service Connect connection propagation allows services in VPC spokes to be transitively and globally accessible from other VPC spokes in the same hub. [Learn more](#)

On
 Off

SAVE CANCEL

□

Deploy GCP NCC VPC Spoke

- Deploy a single GCP NCC Spoke in the host VPC.

Network Connectivity VPC Spoke Creation:

← Add spokes to hub

Project ID
mgillespie-01

Hub name
ncc_hub

Spokes

^ New spoke

Spoke type

VPC network [?]

Producer VPC network [?]

VPN tunnel [?]

VLAN attachment [?]

Router appliance [?]

Spoke name *
vpc-spoke-west

Description

Add VPC network to the spoke

A VPC network is a virtual version of a physical network, implemented inside of Google's production network.

Associated VPC network

VPC network *
gcp-host-project-west-regional

VPC spoke filter (optional) ^

By default, when you associate a VPC network with a spoke, all of its subnets are advertised. However, you can use filters to customize how routes are advertised. [Learn more](#)

Filter actions

IPv4 ranges

Include export private IPv4 subnet ranges from spoke to hub
The default is all private IPv4 ranges. To limit permitted IPv4 ranges, specify the ranges to export.

All private IPv4 subnet ranges (default)

Specify private IPv4 subnet ranges

Exclude export subnet ranges from spoke to hub
Specify the exclude filter to prevent specific IP ranges from being propagated to the hub.

IPv6 ranges

Include export all IPv6 subnet ranges from spoke to hub
IPv6 ranges are not exported by default. To export IPv6 ranges, select this checkbox.

ⁱ After you create a spoke, you can't change IPv4 filters. However, you can change IPv6 filters.

DONE

Deploy GCP NCC Router Appliance Spokes

- Deploy a single GCP NCC Spoke in the Aviatrix BGPoLAN VPC for each region
- Associate the related Aviatrix Transit gateway with the NCC spoke.
- Import all IP ranges from hub to spoke.


Network Connectivity Center Spoke Creation:

← Add spokes to hub


Project ID
mgillespie-01


Hub name
ncc_hub


Spokes


^ New spoke 


Spoke type

VPC network 

Producer VPC network 


VPN tunnel 


VLAN attachment 

Router appliance 

Spoke name *
west-spoke


Description

Region *
us-west1 (Oregon) 

Site-to-site data transfer
Site-to-site data transfer capability is supported only in certain regions. [Learn more](#) 

On

Off

VPC network *
avx-transit-west1-bgpolan 

All resources in spokes with site-to-site data transfer enabled need to come from one VPC network.

Network Connectivity Center Spoke Creation – Adding router appliances:

Add instances to the spoke

Router appliance instances enable connectivity to a VPC network and permit data transfer between on-premises sites. [Create instance](#)

Instance 1

avx-transit-west

Instance 2

avx-transit-west-1

[+ ADD INSTANCE](#)

Hybrid spoke filter (optional)

By default, subnets advertised to the hub are not propagated to hybrid spokes. You can use filters to customize how routes are advertised. [Learn more](#)

[Hub routes](#)

Filter actions

Include import all IPv4 ranges from hub to spoke

This is in addition to the custom route advertisement through Cloud Router. [Learn more](#)

DONE

Deploy GCP Cloud Routers

- Deploy a single GCP Cloud Router in the same Aviatrix BGPoLAN VPC for each region.
- The Cloud Router must be configured with two interfaces. *Aviatrix does not support peering both primary and HA Transits to the same address as of 8.0.0.*
 - This cannot be done via the UI and must be done via Terraform or CLI.

While you can create the initial Cloud Router interfaces via the Google Cloud Console *during the creation of a BGP peer*, the console doesn't directly expose the ability to define one interface as the redundant interface of another in the same way that the `gcloud` command-line tool does with the `--redundant-interface` flag.

To create a truly redundant interface, explicitly linked to a primary interface for high availability, you'll want to use the Google Cloud CLI.

```
gcloud compute routers add-interface [ROUTER_NAME] \  
  --region [REGION] \  
  --interface-name [INTERFACE_NAME] \  
  --ip-address [INTERFACE_IP] \  
  --subnet [SUBNET_NAME] \  
  --peer-ip [PEER_IP] \  
  --redundant-interface
```

- The Cloud Router must be configured to advertise “all subnets visible” in addition to the Google API ranges (199.36.153.4/30 and 199.36.153.8/30) to the Aviatrix Transit.

Google Cloud Router Configuration:

← Create a cloud router

Google Cloud Router dynamically exchanges routes between your Virtual Private Cloud (VPC) and on-premises networks by using Border Gateway Protocol (BGP).

| | | |
|---|---------------------------|-----------|
| Name * | avx-transit-central1-cr | ? |
| <small>Lowercase letters, numbers, hyphens allowed</small> | | |
| Description | | |
| Network * | avx-transit-central-bgp-0 | ? |
| Region * | us-central1 (Iowa) | ? |
| Cloud Router ASN | 65202 | ? |
| BGP peer keepalive interval | | seconds ? |
| BGP identifier | | ? |
| <small>E.g. 169.254.16.16/30. If not specified, Google will automatically assign an IPv4 range.</small> | | |

Advertised routes ?

Routes

- Advertise all subnets visible to the Cloud Router (Default)
- Create custom routes

Advertise all subnets

- Advertise all subnets visible to the Cloud Router

Filter ?

| Subnet ↑ | IP ranges |
|---------------------------------------|--------------------|
| avx-transit-central-bgp-0-us-central1 | IPv4 : 10.2.0.0/27 |

Google Cloud Router Configuration – Custom Advertisement:

Advertised routes

Routes

- Advertise all subnets visible to the Cloud Router (Default)
- Create custom routes

Advertise all subnets

- Advertise all subnets visible to the Cloud Router

 **Filter** 

| Subnet  | IP ranges |
|--|--------------------|
| avx-transit-central-bgp-0-us-central1 | IPv4 : 10.2.0.0/27 |

Custom ranges

Add IPv4 and IPv6 ranges to advertise

| | |
|---|---|
|  199.38.153.4/30 |  |
|  199.38.153.8/30 |  |

[ADD A CUSTOM ROUTE](#)

Configure VPC Static Routes

- Configure static routes in the BGPoLAN VPC to 199.36.154.4/30 and 199.36.154.8/30

Configure BGP Peerings

- Create BGP peerings between the GCP Cloud Router and Aviatrix Transit Gateways.
- Cloud Router interface 1 should peer with the Primary Aviatrix Transit.
- Cloud Router interface 2 should peer with the HA Aviatrix Transit.

Aviatrix Transit Gateway BGP:

| Name | Transit Type | Local Gateway | Local Gateway | Remote Gateway | Local Subnet | Remote Subnet | Local ASN | Remote ASN | BGP Local IP | BGP Neighbor IP | BGP Neighbor Status |
|-----------------------------------|--------------|---------------|---------------|-------------------------|--------------|---------------|-----------|------------|--------------|-----------------|---------------------|
| avx-transit-central-cr-ncc-to-avx | BGP over LAN | | | avx-transit-central-tag | 65301 | 65302 | 10.2.0.2 | 10.2.0.9 | Established | | |
| avx-transit-east-cr-ncc-to-avx | BGP over LAN | | | avx-transit-east-tag | 65262 | 65262 | 10.5.0.2 | 10.5.0.8 | Established | | |

Cloud Router BGP:

| Instance Name | IP address | Cloud router | BGP session name | Cloud router BGP IP | BGP peer IP |
|--------------------------|------------|------------------------|--|---------------------|-------------|
| avx-transit-central | 10.2.0.2 | avx-transit-central-cr | avx-transit-central-cr-pri-to-avx-transit-central-avx-pri-peer | 10.2.0.9 | 10.2.0.2 |
| CONFIGURE BGP SESSION | | | | | |
| avx-transit-central-hagw | 10.2.0.3 | avx-transit-central-cr | avx-transit-central-cr-pri-to-avx-transit-central-avx-ha-peer | 10.2.0.8 | 10.2.0.3 |
| CONFIGURE BGP SESSION | | | | | |

Configure Aviatrix Firenet Inspection

- Enable Firenet inspection of S2C BGPoLAN.

Aviatrix Firenet Inspection:



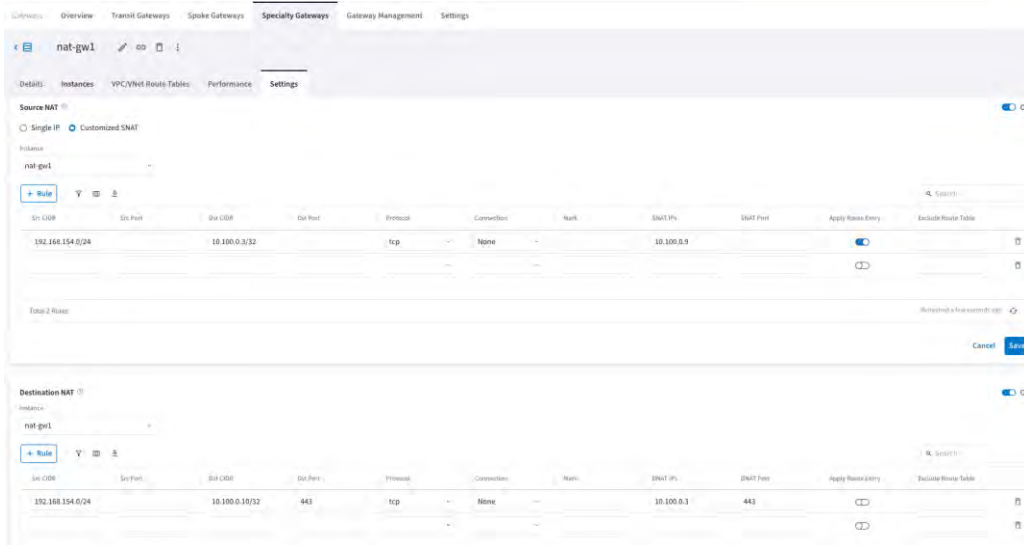
Deploy Google Kubernetes Engine

- Deploy GKE in the appropriate shared VPC subnet.

Deploy an Aviatrix Gateway for SNAT/DNAT.

- Deploy an Aviatrix Gateway in the shared VPC subnet. For resilience, it is best practice to install a gateway in every zone for which GKE is configured.
- To deploy the gateways in multiple regions, ensure the gateway is launched as a specialty gateway and NOT a spoke gateway.
- Configure DNAT rules to accept traffic from an internal GCP load balancer and redirect to the GKE endpoint.
- Configure SNAT rules to source NAT traffic from Aviatrix Gateway to GKE endpoint to the address of the Aviatrix Gateway.
- Note: When upgrading an Aviatrix gateway that's associated to a load-balancer, you ***must reassociate the gateway to the load balancer afterwards.***

Aviatrix Gateway – SNAT/DNAT Configuration:



Aviatrix Gateway – SNAT/DNAT Configuration Summary:

| DNAT Policy | | SNAT Policy | |
|------------------|-----------------------|------------------|-----------------------|
| Source CIDR | On-premise addresses | Source CIDR | On-premise addresses |
| Destination CIDR | Load Balancer Address | Destination CIDR | GKE Endpoint |
| Protocol | TCP | Protocol | TCP |
| Interface | Load Balancer Address | Interface | Load Balancer Address |
| Connection | On-premise addresses | Connection | On-premise addresses |
| DNAT IPs | GKE Endpoint | SNAT IPs | Aviatrix Gateway IP |
| DNAT Port | 443 | | |

Configure GCP Firewall rules for Health Checks.

- Configure GCP Firewall rules to allow access from 35.191.0.0/16 and 130.211.0.0/22 to the Load Balancer backend pool (Aviatrix Gateways). This can be simplified as allowing access to the entire VPC is desired.

Example health check firewall rule:

The screenshot shows the configuration for a GCP Firewall rule named 'health-checks'. At the top, there are navigation links: a back arrow, 'Firewall rule details', an 'Edit' button with a pencil icon, and a 'Delete' button with a trash icon. The rule name 'health-checks' is displayed at the top left. Below it, the 'Logs' section is set to 'On', with a link to 'view in Logs Explorer' and a 'Show logs details' button. The 'Network' is 'gcp-host-project'. The 'Priority' is '15'. The 'Direction' is 'Ingress'. The 'Action on match' is 'Allow'. The 'Tags' section is empty. The 'Source filters' section contains two IP ranges: '35.191.0.0/16' and '130.211.0.0/22'. The 'Destination filters' section contains one IP range: '10.100.0.0/16'. The 'Protocols and ports' are set to 'All'. The 'Enforcement' is 'Enabled'. The 'Insights' are set to 'None'.

health-checks

Logs ⓘ
On
[view in Logs Explorer](#)
[Show logs details](#)

Network
gcp-host-project

Priority
15

Direction
Ingress

Action on match
Allow

Tags
-

Source filters

| | |
|-----------|----------------|
| IP ranges | 35.191.0.0/16 |
| | 130.211.0.0/22 |

Destination filters

| | |
|-----------|---------------|
| IP ranges | 10.100.0.0/16 |
|-----------|---------------|

Protocols and ports
All

Enforcement
Enabled

Insights
None

Deploy a GCP Load Balancer

- Deploy a GCP Internal network load balancer in passthrough mode.
- Configure the front-end to listen on TCP 443
- Configure target groups to include the Aviatrix Gateway
- Configure health checks to test the response of TCP:443
- Assign the newly created target group as a backend.
 - Note: When performing Aviatrix Image upgrades to the NAT Spoke, the instance is replaced. As a result, ***ensure you update the backend group to include the newly created instance.***

GCP Internal Load Balancer:

The screenshot displays the configuration for a GCP Internal passthrough Network Load Balancer named 'api-endpoint-lb'. The configuration is divided into Frontend and Backend sections.

Frontend Configuration:

| Protocol | IP version | Scope | Subnetwork | IP:Ports | DNS name |
|----------|------------|--------------------------------|----------------------|-----------------|----------|
| TCP | IPv4 | us-central1 with global access | subnet-us-central1-0 | 10.100.0.10:443 | |

Backend Configuration:

| Region | Network | Endpoint protocol | Session affinity | Health check | Logging |
|-------------|------------------|-------------------|------------------------------|------------------|----------|
| us-central1 | gcp-host-project | TCP | Client IP, port and protocol | tcp-health-check | Disabled |

Advanced Configurations:

| Instance group | Type | IP stack type | Scope | Healthy | Autoscaling | Use as failover group |
|--------------------|----------------|---------------|---------------|---------|------------------|-----------------------|
| unmanaged-vm-group | Instance group | IPv4 | us-central1-a | 1 of 1 | No configuration | No |

GCP Internal Load Balancer – Backend Instance Group:

← unmanaged-vm-group [Edit](#) [Delete Group](#)

[Overview](#) [Details](#) [Monitoring](#) [Errors](#)

Instances by status

1 instance ?

✓ 1

Network

gcp-host-project

Status Unmanaged

Creation Time Apr 2, 2025, 11:10:15 AM UTC-04:00

Description

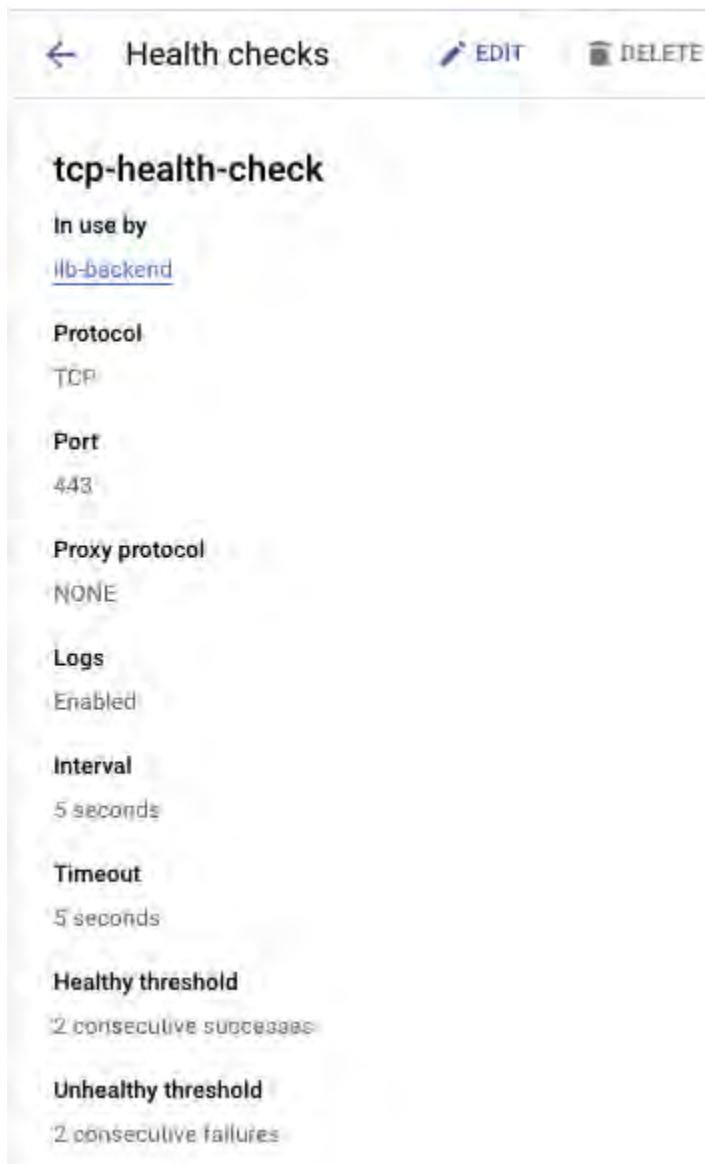
Location us-central1-a

VM instances [Suspend](#) [Stop](#) [Start / Resume](#) [Remove from group](#) [Delete](#)

Filter Enter property name or value

| <input type="checkbox"/> Status | Name ↑ | Creation Time | Template | Per instance config | Internal IP | External IP | Health Check Status |
|---|-------------------------|------------------------------------|----------|---------------------|-------------------------------------|--------------|---------------------|
| <input type="checkbox"/> ✓ | nat-gw1 | Apr 2, 2025, 10:24:06 AM UTC-04:00 | - | | 10.100.0.9 (nic0) | 34.56.60.140 | |

GCP Internal Load Balancer – Health Check:



The screenshot shows the configuration for a health check named 'tcp-health-check'. At the top, there is a navigation bar with a back arrow, the title 'Health checks', and 'EDIT' and 'DELETE' buttons. The configuration details are as follows:

- tcp-health-check**
- In use by:** [ilb-backend](#)
- Protocol:** TCP
- Port:** 443
- Proxy protocol:** NONE
- Logs:** Enabled
- Interval:** 5 seconds
- Timeout:** 5 seconds
- Healthy threshold:** 2 consecutive successes
- Unhealthy threshold:** 2 consecutive failures

Establish DNS

- Deploy a GCP Internal network load balancer in passthrough mode.

Example NCC spoke configuration:

^ New spoke 🗑

Spoke type

VPC network ?

Producer VPC network ?

VPN tunnel ?

VLAN attachment ?

Router appliance ?

Spoke name *

Description

Region * ?

Site-to-site data transfer

Site-to-site data transfer capability is supported only in certain regions. [Learn more](#) 🔗

On

Off

VPC network *

All resources in spokes with site-to-site data transfer enabled need to come from one VPC network.

Add instances to the spoke

Router appliance instances enable connectivity to a VPC network and permit data transfer between on-premises sites. [Create instance](#)

Instance 1

[+ ADD INSTANCE](#)

[+ ADD INSTANCE](#)

Hybrid spoke filter (optional)

By default, subnets advertised to the hub are not propagated to hybrid spokes. You can use filters to customize how routes are advertised. [Learn more](#) 🔗

[Hub routes](#)

Filter actions

Include import all IPv4 ranges from hub to spoke

This is in addition to the custom route advertisement through Cloud Router. [Learn more](#) 🔗

[DONE](#)

Using round-robin DNS vs. GCP Load Balancer (optional)

- To simplify management, round-robin DNS can be configured as an alternative to a load balancer.

Design Validation

After implementing any solution, thorough testing is needed to ensure correct design implementation and requirement fulfillment.

Monitoring and Management Best Practices

Aviatrix CoPilot alerts should be configured to ensure proper operation of the solution. The out-of-the-box alerts serve as a guideline for monitoring overall health. Additional alerts should be configured to detect the following:

- High CPU and Memory usage on Aviatrix Transit Gateways.
- Packet drops.
- BGP peering issues.

Example Global Network Health Configuration:

Edit Alert Configuration: Global Network Health

Name
Global Network Health System Defined Alert

Monitor

- Controller
- CoPilot
- Gateways

All Gateways Selected

Condition

Matches any condition (OR)

Gateway Status Down Keep Alive Fail Config Fail Upgrade Fail

| | | |
|-----------------------------|-----------|----|
| PPS Limit Exceeded Drop (%) | more than | 1 |
| Packet Drop (%) | more than | 5 |
| Memory Used (%) | more than | 90 |

Monitor network metrics on selected interfaces (individually) Off

Evaluation Period: 15 min
Minimum Count of Matching Entities: 1

Send Alerts To

Recipients: network-operations-center

Sends separate alerts for each gateway On

[Reset to Defaults](#) [Cancel](#) [Save](#)

Example High CPU Configuration:

Create Alert Configuration

Name
Transit Gateway - High CPU

Monitor

Controller

CoPilot

Gateways

avx-transit-central x avx-transit-central-hagw x avx-transit-east x avx-transit-east-hagw x

Condition

Matches all conditions (AND) v

CPU Used (%) x v more than v 80

Evaluation Period: 10 min Minimum Count of Matching Entities: 1

Send Alerts To

Recipients

network-operations-center x

Send separate alerts for each gateway On

Cancel Save

Example Memory Configuration:

Edit Alert Configuration: Global Memory Swap Surge

Name
Global Memory Swap Surge System Defined Alert

Monitor

Controller

CoPilot

Gateways

All Gateways Selected

Condition

Matches all conditions (AND) ▼

| | | | |
|----------------|-----------|-----------|---|
| Memory Swapped | more than | 0 | B |
| Memory Total | more than | 107374182 | |

Evaluation Period ⓘ 15 min

Minimum Count of Matching Entities ⓘ 1

Send Alerts To

Recipients

network-operations-center x

Send separate alerts for each gateway On

[Reset to Defaults](#) [Cancel](#) [Save](#)

Additionally, [alerting policies should be configured in the GCP console](#) to detect similar issues with the Google Cloud Routers and Load Balancers.

Firewalls should be managed and monitored by the appropriate software.

