Cloud VPN overview

This page describes concepts related to Google Cloud VPN.

Cloud VPN securely connects your peer (/vpn/docs/concepts/key-terms#peer-definition) network to your Virtual Private Cloud (VPC) network (/vpc/docs) through an IPsec (https://wikipedia.org/wiki/IPsec) VPN (https://wikipedia.org/wiki/Virtual_private_network) connection. Traffic traveling between the two networks is encrypted by one VPN gateway, and then decrypted by the other VPN gateway. This protects your data as it travels over the internet. You can also connect two instances of Cloud VPN to each other.

To create a virtual private network (VPN), see Choosing a VPN option (/vpn/docs/how-to/choosing-a-vpn).

For definitions of terms used in Cloud VPN documentation, see Key terms (/vpn/docs/concepts/key-terms).

Choosing a hybrid networking solution

To determine whether to use Cloud VPN, Dedicated Interconnect, or Partner Interconnect as your hybrid networking connection to Google Cloud, see the following resources:

- Cloud VPN and other hybrid connectivity solutions (/vpn/docs/concepts/choosing-a-hybrid-solution)
- How to choose a Cloud Interconnect connection type (/interconnect/docs/how-to/choose-type)

Types of Cloud VPN

Google Cloud offers two types of Cloud VPN gateways, HA VPN and Classic VPN.

For information about moving to HA VPN, see Moving to HA VPN from Classic VPN (/vpn/docs/how-to/moving-to-ha-vpn).

HA VPN

HA VPN is a high-availability (HA) Cloud VPN solution that lets you securely connect your on-premises network to your Virtual Private Cloud network through an IPsec VPN connection in single region. HA VPN provides an SLA of 99.99% service availability.

When you create an HA VPN gateway, Google Cloud automatically chooses two public IP addresses, one for each of its fixed number of two interfaces. Each IP address is automatically chosen from a unique address pool to support high availability. Each of the HA VPN gateway interfaces supports multiple tunnels. You can also create multiple HA VPN gateways.

You can configure an HA VPN gateway with only one active interface and one public IP address; however, *this configuration does not provide a 99.99% service availability SLA*.

HA VPN gateways are referred to as *VPN gateways*, rather than *target VPN gateways*, in the API documentation and in gcloud commands.

You don't need to create any forwarding rules for HA VPN gateways.

HA VPN uses an external VPN gateway resource in Google Cloud to provide information to Google Cloud about your peer VPN gateway or gateways. For more information, see the definitions for external VPN gateway resource (/vpn/docs/concepts/key-terms#external-vpn-gateway-definition) and peer VPN gateway (/vpn/docs/concepts/key-terms#peer-definition).

The following diagram shows the HA VPN concept, showing a topology that includes the two interfaces of a HA VPN gateway connected to two peer VPN gateways. For more detailed HA VPN topologies (configuration scenarios), see the Cloud VPN Topologies page. (/vpn/docs/concepts/topologies)

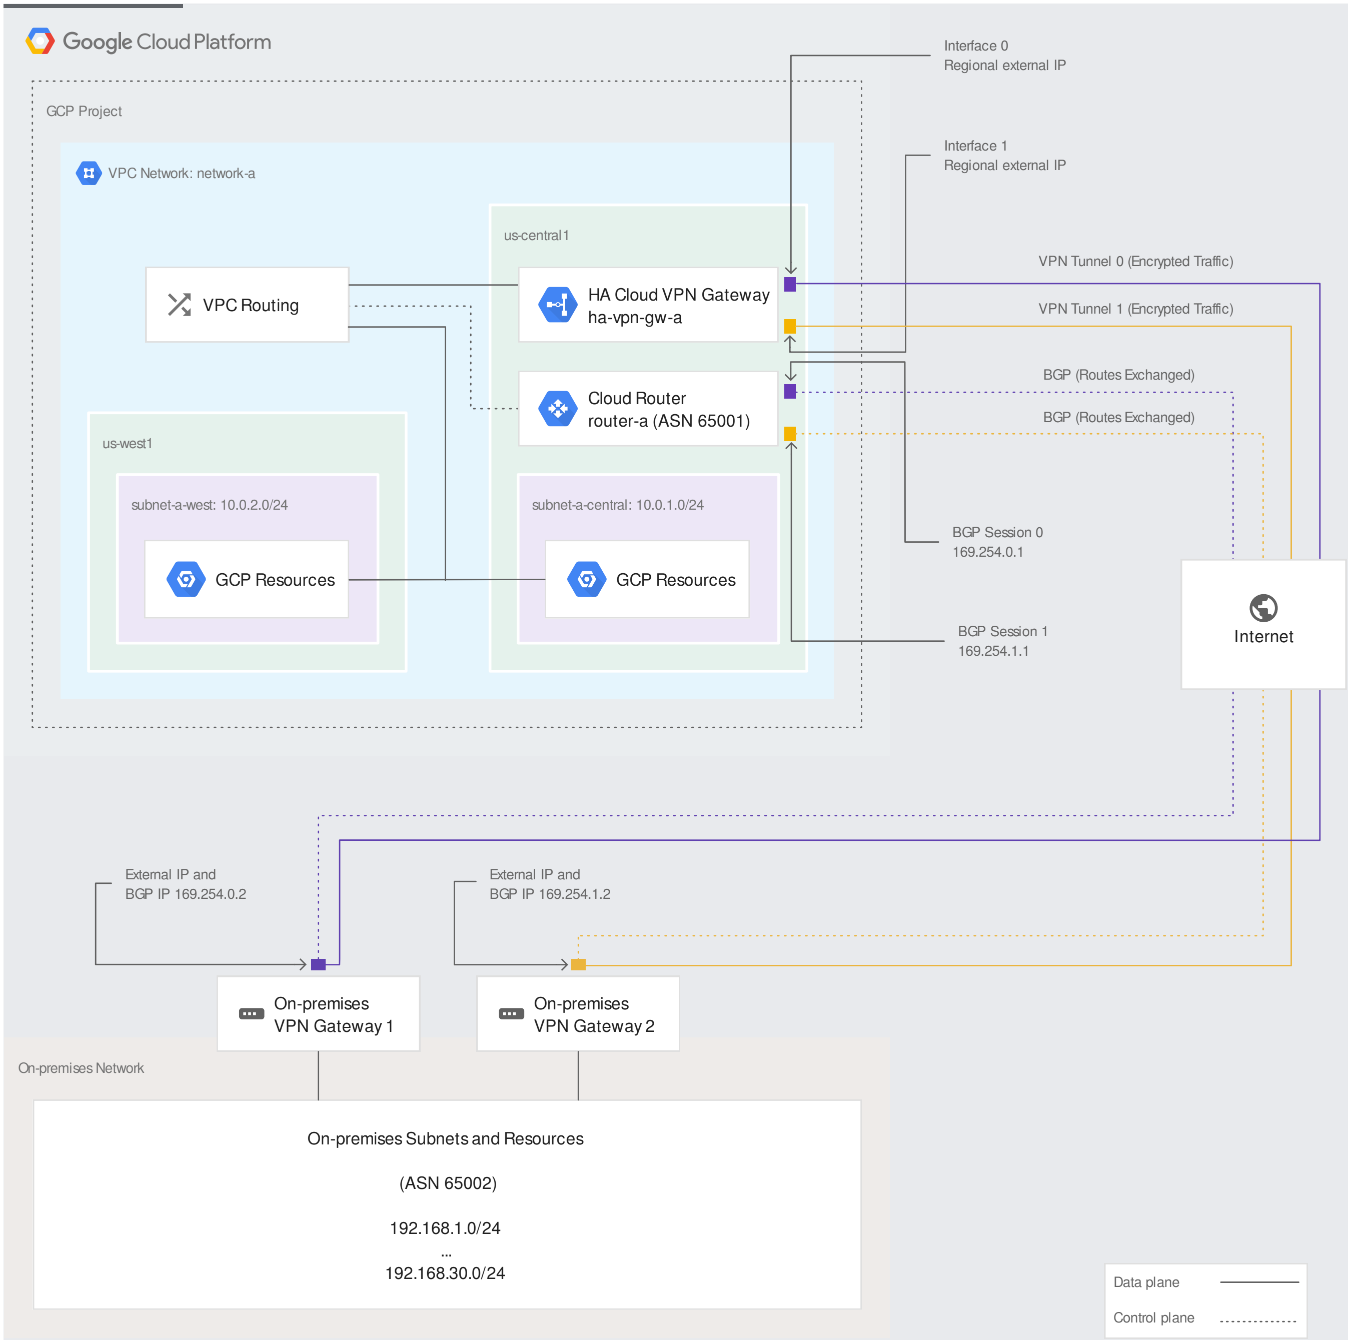

(/vpn/images/ha-vpn-gcp-to-on-prem-2-a.svg)
A HA VPN gateway to two peer VPN gateways (click to enlarge)

HA VPN requirements

Your Cloud VPN configuration must meet the following requirements to achieve a service-level availability of 99.99% for HA VPN:

- When you connect an HA VPN gateway to your peer gateway, 99.99% availability is guaranteed only on the Google Cloud side of the connection. End-to-end availability is subject to proper configuration of the peer VPN gateway.

- If both sides are Google Cloud gateways and are properly configured, end-to-end 99.99% availability is guaranteed.
- To achieve high availability when both VPN gateways are located in VPC networks, you must use two HA VPN gateways, and both of them must be located in the same region. Even though both gateways must be located in the same region, the routes to their subnets that they share with each other can be located in any region if your Virtual Private Cloud network uses *global dynamic routing mode*. If your VPC network uses *regional dynamic routing mode*, only routes to subnets in the same region are shared with the peer network, and learned routes are applied only to subnets in the same region as the VPN tunnel. For more information about the dynamic routing mode of a VPC network, refer to the VPC networks overview. (/vpc/docs/vpc#routing_for_hybrid_networks)
- HA VPN rejects Google Cloud IP addresses when they are configured in an external VPN gateway resource. An example of this is using the external IP address of a VM instance as the public IP address for the external VPN gateway resource. The only supported HA VPN Google Cloud-to-Google Cloud topology is where HA VPN is used on both sides, as documented in Creating Google Cloud-to-Google Cloud HA VPN gateways. (/vpn/docs/how-to/creating-ha-vpn2)
- You must configure two VPN tunnels from the perspective of the Cloud VPN gateway:
  - If you have *two peer VPN gateway devices*, each of the tunnels from each interface on the Cloud VPN gateway must be connected to its own peer gateway.
  - If you have *a single peer VPN gateway device with two interfaces*, each of the tunnels from each interface on the Cloud VPN gateway must be connected to its own interface on the peer gateway.
  - If you have *a single peer VPN gateway device with a single interface*, both of the tunnels from each interface on the Cloud VPN gateway must be connected to the same interface on the peer gateway.
- A peer VPN device must be configured with adequate redundancy. The details of an adequately redundant configuration are specified by the device vendor, and may or may not include multiple hardware instances. Refer to the vendor documentation for the peer VPN device for details. If two peer devices are required, each peer device must be connected to a different HA VPN gateway interface. If the peer side is another cloud provider like AWS, VPN connections must be configured with adequate redundancy on the AWS side as well.
- Your peer VPN gateway device must support dynamic (BGP) routing.

Classic VPN

All Cloud VPN gateways created before the introduction of HA VPN are considered Classic VPN gateways. To move f
c VPN to HA VPN, see the detailed instructions. (/vpn/docs/how-to/moving-to-ha-vpn)

In contrast, Classic VPN gateways have a single interface, a single external IP address, and support tunnels using dynamic (BGP) (/vpn/docs/concepts/key-terms#bgp-definition) or static routing (route based or policy based). They provide an SLA of 99.9% service availability.

For supported Classic VPN topologies, see the Classic VPN topologies page. (/vpn/docs/concepts/classic-topologies)

Classic VPNs are referred to as *target VPN gateways* in the API documentation and in gcloud commands.

Comparison table

The following table compares HA VPN features with those for Classic VPN.

The tunnel API resource and tunnel configuration remain the same for both Classic VPN and HA VPN.

Feature	HA VPN	Classic VPN
SLA	Provides a 99.99% SLA when configured with two interfaces and two public IPs	Provides a 99.9% SLA
Creation of public IPs and forwarding rules	Public IPs created from a pool. No forwarding rules required	Public IPs and forwarding rules must be created
Routing options supported	Only Dynamic Routing (BGP)	Static Routing (policy based, route based) or Dynamic Routing using BGP
Two tunnels from one Cloud VPN gateway to the same peer gateway	Supported	Not supported
API resources	Known as the vpn-gateway resource	Known as the target-vpn-gateway resource

Specifications

Cloud VPN has the following specifications:

- Cloud VPN can be used with VPC networks and legacy networks. For VPC, custom mode is recommended so that you have full control over the ranges of IP addresses used by the subnets in the network. For more information, see the documentation for VPC networks in general (/vpc/docs/vpc), legacy networks (/vpc/docs/legacy), and custom mode networks (/vpc/docs/vpc#subnet-ranges).
    - Classic VPN and HA VPN gateways use external (internet routable) IPv4 addresses. Only ESP, UDP 500, and UDP 4500 traffic is permitted to these addresses. This applies to Cloud VPN addresses configured by you for Classic VPN or to automatically assigned addresses for HA VPN.
    - If IP address ranges for on-premises subnets overlap with IP addresses used by subnets in your VPC network, see Order of routes (/vpn/docs/concepts/order-of-routes) to determine how routing conflicts are resolved.
- Cloud VPN can be used in conjunction with *Private Google Access for on-premises hosts*. For more information, see Private Google Access options (/vpc/docs/private-access-options).
- Each Cloud VPN gateway must be connected to another Cloud VPN gateway or a peer VPN gateway.
- The peer VPN gateway must have a static external (internet routable) IPv4 address. You need to know its IP address in order to configure Cloud VPN.
    - If your peer VPN gateway is behind a firewall, you must configure the firewall to pass ESP (IPsec) protocol and IKE (UDP 500 and UDP 4500) traffic to it. If the firewall provides Network Address Translation (NAT), see UDP encapsulation and NAT-T (#udp_and_nat).
- Cloud VPN requires that the peer VPN gateway be configured to support prefragmentation. Packets must be fragmented *before* being encapsulated.
- Cloud VPN uses replay detection with a window of 4096 packets. You cannot turn this off.

Network bandwidth

Each Cloud VPN tunnel can support up to 3 Gbps. Actual bandwidth depends on several factors:

- The network connection between the Cloud VPN gateway and your peer gateway:
    - **Network bandwidth between the two gateways.** Throughput is higher if you have established a Direct Peering (/interconnect/direct-peering) relationship with Google than if your VPN traffic is sent over the public internet.
    - **Round Trip Time (RTT)** (https://wikipedia.org/wiki/Round-trip_delay_time) **and packet loss.** Elevated RTT and/or packet loss rates greatly reduce TCP performance.

- **Capabilities of your peer VPN gateway.** See your device's documentation for more information.
- **Packet size.** Cloud VPN uses a Maximum Transmission Unit (MTU) (https://wikipedia.org/wiki/Maximum_transmission_unit) of 1460 bytes. Peer VPN gateways must be configured to use an MTU of no greater than 1460 bytes. Because processing happens on a per-packet basis, for a given packet rate, a significant number of smaller packets can reduce overall throughput. To account for ESP overhead, you might also need to set the MTU values for *systems* sending traffic through VPN tunnels to values less than the MTU of the tunnel. See MTU considerations (/vpn/docs/concepts/mtu-considerations) for a detailed discussion and recommendations.
- **Packet rate.** For ingress and egress, the recommended maximum packet rate for each Cloud VPN tunnel is 250,000 packets per second (pps). If you need to send packets at a higher rate, you must create more VPN tunnels.

When measuring TCP bandwidth of a VPN tunnel, you should measure more than one simultaneous TCP stream. If you are using the *iperf tool* (https://wikipedia.org/wiki/Iperf), use the `-P` parameter to specify the number of simultaneous streams.

IPsec and IKE support

- Cloud VPN supports IKEv1 (https://tools.ietf.org/html/rfc2409) and IKEv2 (https://tools.ietf.org/html/rfc5996) using a shared secret (IKE pre-shared key) (/vpn/docs/how-to/generating-pre-shared-key) and these IKE ciphers (/vpn/docs/concepts/supported-ike-ciphers).
- Cloud VPN supports ESP in tunnel mode (https://wikipedia.org/wiki/IPsec#Tunnel_mode) with authentication, but does not support AH (https://wikipedia.org/wiki/IPsec#Authentication_Header) or ESP in transport mode (https://wikipedia.org/wiki/IPsec#Transport_mode).

Note that Cloud VPN does not perform policy-related filtering on incoming authentication packets. Outgoing packets are filtered based on the IP range configured on the Cloud VPN gateway.

- Cloud VPN only supports a pre-shared key (shared secret) for authentication. You must specify a shared secret when you create the Cloud VPN tunnel. This same secret must be specified when creating the tunnel at the peer gateway. See these guidelines for creating a strong shared secret (/vpn/docs/how-to/generating-pre-shared-key).
- For ciphers and configuration parameters supported by Cloud VPN, see Supported IKE ciphers (/vpn/docs/concepts/supported-ike-ciphers).

UDP encapsulation and NAT-T

For information about how to configure your peer device to support NAT-T with Cloud VPN, see the UDP and NAT-T section in the Advanced overview (/vpn/docs/concepts/advanced#udp-encapsulation).

Cloud VPN as a transit network

Carefully review the Google Cloud Service specific terms (/terms/service-terms-20170718#15.-additional-restrictions) before you use Cloud VPN.

Do not use Cloud VPN tunnels to connect two or more on-premises networks for the sole purpose of passing traffic through a VPC network as a transit network. Hub-and-spoke configurations like this are a violation of the Google Cloud Service Specific Terms.

Active/active and active/passive routing options for HA VPN

If a Cloud VPN tunnel goes down, it restarts automatically. If an entire virtual VPN device fails, Cloud VPN automatically instantiates a new one with the same configuration. The new gateway and tunnel connect automatically.

VPN tunnels connected to HA VPN gateways must use dynamic (BGP) routing. Depending on the way you configure route priorities for HA VPN tunnels, you can create an active/active or active/passive routing configuration. For both of these routing configurations, both VPN tunnels remain active.

The following table compares the features of an active/active or active/passive routing configuration.

Feature	Active/active	Active/passive
Throughput	The effective aggregate throughput is the *combined throughput of both tunnels*.	After reducing from two active tunnels to one, *the effective overall throughput is cut in half*, which can result in slower connectivity or dropped packets.

Feature	Active/active	Active/passive
Route advertisement	Your peer gateway advertises the peer network's routes with *identical MED values for each tunnel*. The Cloud Router managing the Cloud VPN tunnels imports these as custom dynamic routes in your VPC network with *identical priorities*. Egress traffic sent to your peer network uses Equal Cost Multi-path (ECMP) routing (https://wikipedia.org/wiki/Equal-cost_multi-path_routing) . The same Cloud Router also advertises routes to your VPC network using *identical priorities*. Your peer gateway can use these routes to send egress traffic to Google Cloud *using ECMP* as well.	Your peer gateway advertises the peer network's routes with *different MED values for each tunnel*. The Cloud Router managing the Cloud VPN tunnels imports these as custom dynamic routes in your VPC network with *different priorities*. Egress traffic sent to your peer network uses the route with the highest priority, as long as the associated tunnel is available. The same Cloud Router also advertises routes to your VPC network using *different priorities for each tunnel*. Your peer gateway can only send traffic to Google Cloud *using the tunnel with highest priority*.
Failover	*if one tunnel becomes unavailable, Cloud Router withdraws the learned custom dynamic routes whose next hops are the unavailable tunnel.* This withdrawal process can take up to 40 seconds (/vpn/docs/concepts/order-of-routes#tunnels_down) , during which packet loss is expected.	*Uses a maximum of one tunnel at a time, so that the second tunnel is able to handle all of your egress bandwidth* in the event that the first tunnel fails and needs to be failed over. If one tunnel becomes unavailable, Cloud Router withdraws the learned custom dynamic routes whose next hops are the unavailable tunnel. This withdrawal process can take up to 40 seconds (/vpn/docs/concepts/order-of-routes#tunnels_down) , during which packet loss is expected.

Using multiple tunnels or gateways

For an example of a multiple-tunnel active/passive scenario, see the Topology page
/docs/concepts/topologies#more-bandwidth).

on: We recommend that you do not use an active/passive configuration when you have more than one HA VPN gatew
e an active/passive configuration across multiple HA VPN gateways, with an active and passive tunnel pair configur
ateway, HA VPN won't use the passive tunnels for failover until all of the active tunnels on all gateways have failed.
uring multiple gateways with an active/passive configuration *can cause bandwidth loss*.

Depending on the peer gateway configuration, it's possible to construct routes such that some traffic would traverse one tunnel and other traffic would traverse another tunnel due to route priorities (MED values). Similarly, you can adjust the base priority that the Cloud Router uses to share your VPC network routes. These situations demonstrate possible routing configurations that are neither purely active/active nor purely active/passive.

Recommended routing option

When using a single HA VPN gateway, we recommend using an active/passive routing configuration. With this configuration, the observed bandwidth capacity at the time of normal tunnel operation matches the bandwidth capacity observed during failover. This type of configuration is easier to manage, since the observed bandwidth limit stays constant, except for the multiple gateway scenario described previously.

When using multiple HA VPN gateways, an active/active configuration is recommended. With this configuration, the observed bandwidth capacity at the time of normal operation is twice that of the guaranteed bandwidth capacity. However, this configuration effectively underprovisions the tunnels and can cause dropped traffic in case of failover.

Maintenance and availability

Cloud VPN undergoes periodic maintenance. During maintenance, Cloud VPN tunnels are taken offline, resulting in brief drops in network traffic. When maintenance completes, Cloud VPN tunnels are automatically re-established.

Maintenance for Cloud VPN is a normal operational task that may happen at any time without prior notice. Maintenance periods are designed to be short enough so that the Cloud VPN SLA (/vpn/sla) is not impacted.

HA VPN is the recommended method of configuring highly available (HA) VPNs. For configuration options, see the HA VPN topologies page (/vpn/docs/concepts/topologies). If you are using Classic VPN for redundancy and high-throughput options, see the Classic VPN topologies page (/vpn/docs/concepts/classic-topologies).

Best practices

Use these best practices (/vpn/docs/resources/best-practices) to build your Cloud VPN in the most effective way.

What's next

- To set up different types of Cloud VPN gateways, see Choosing a VPN option (/vpn/docs/how-to/choosing-a-vpn).
- To learn about high-availability, high-throughput scenarios or multiple subnet scenarios, see Advanced configurations (/vpn/docs/concepts/advanced).
- To create a custom VPC network, see Creating a custom mode network (/vpc/docs/using-vpc#create-custom-network).
- To maintain VPN tunnels and gateways, see Maintaining VPNs (/vpn/docs/how-to/maintaining-vpns).
- To view Cloud Logging and Cloud Monitoring metrics, see Viewing logs and metrics (/vpn/docs/how-to/viewing-logs-metrics).
- To monitor and solve common issues with Cloud VPN, see Troubleshooting (/vpn/docs/support/troubleshooting).

Next
Cloud VPN and other hybrid connectivity solutions →
(/vpn/docs/concepts/choosing-a-hybrid-solution)

Last updated 2020-04-13.