Cloud VPN (https://cloud.google.com/vpn/)
Documentation (https://cloud.google.com/vpn/docs/) Guides

Choosing a VPN option

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

For information on moving to HA VPN, see Moving to HA VPN from Classic VPN
(https://cloud.google.com/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 Google Cloud 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.

In addition, you don't need to create any forwarding rules for HA VPN gateways.

HA VPN uses an external VPN gateway resource
(https://cloud.google.com/vpn/docs/concepts/overview#external-vpn-gateway-definition) in Google Cloud to provide information to Google Cloud about your peer VPN gateway or gateways
(https://cloud.google.com/vpn/docs/concepts/overview#peer-definition).

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.**

The following diagram shows the HA VPN concept. For more detailed HA VPN topologies (configuration scenarios), see the Cloud VPN Topologies page.

(https://cloud.google.com/vpn/docs/concepts/topologies)

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

The diagram below shows the two interfaces of a HA VPN gateway connected to two peer VPN gateways:
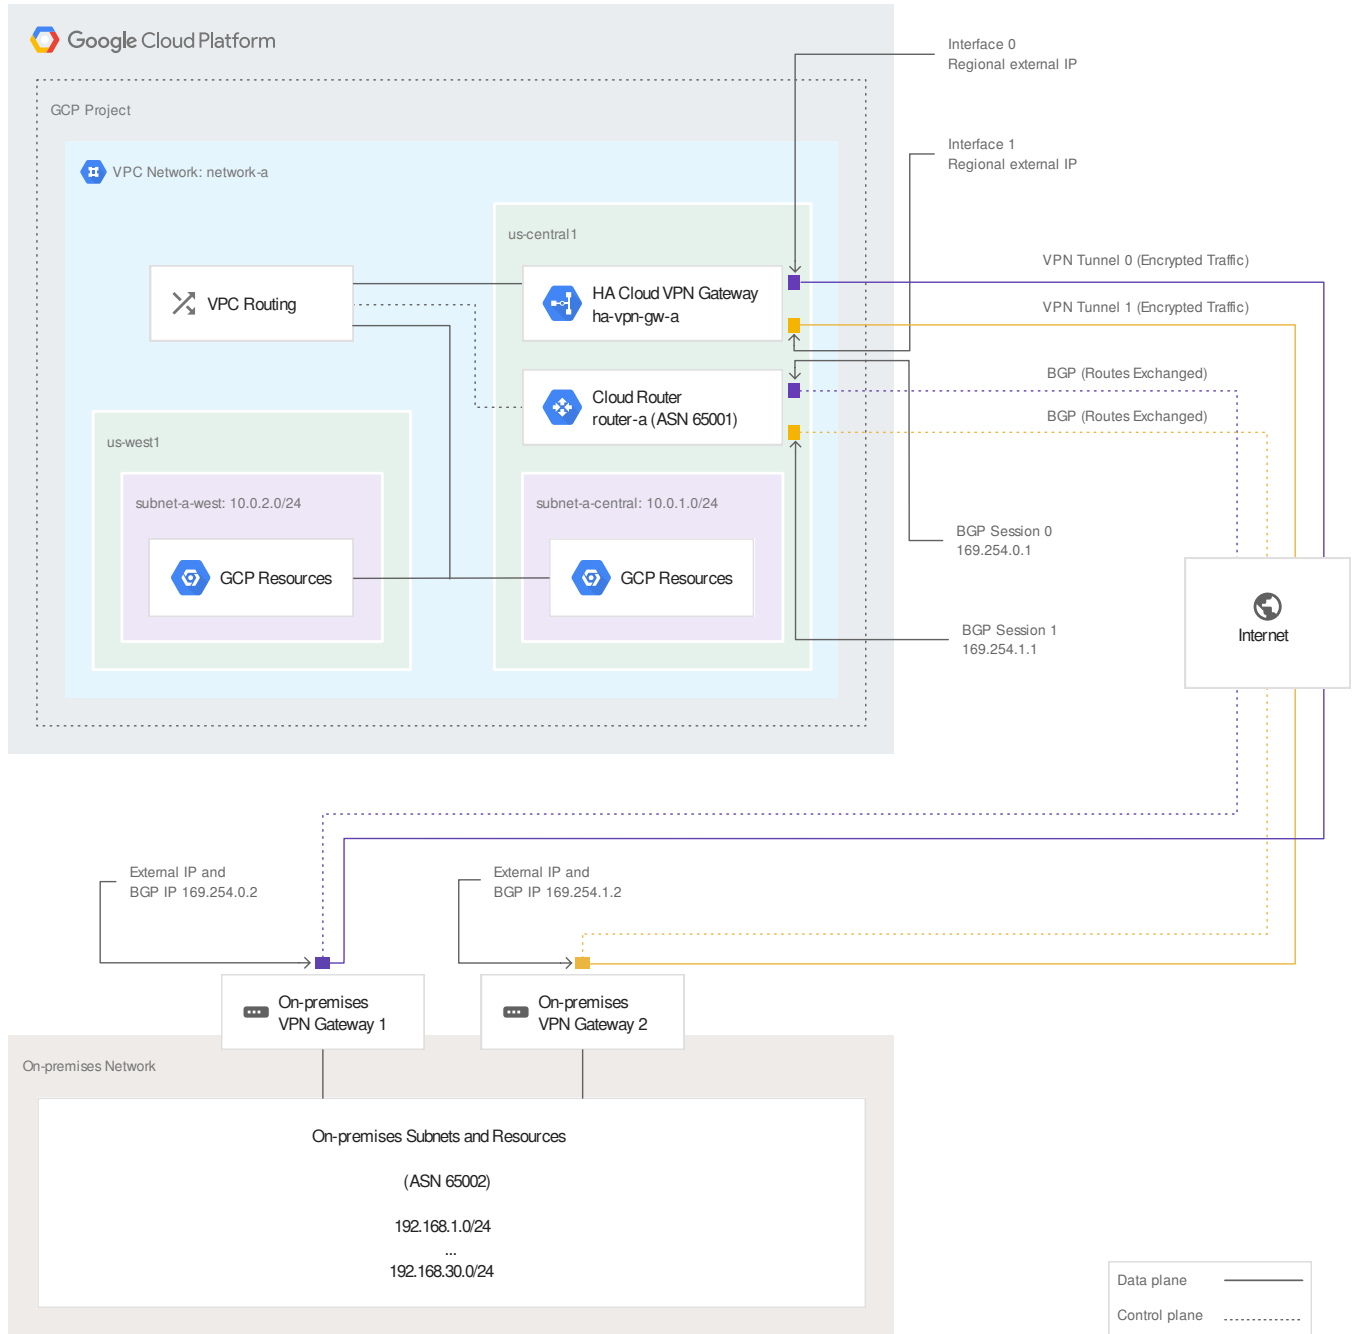
(https://cloud.google.com/vpn/images/ha-vpn-gcp-to-on-prem-2-a.svg)
A HA VPN gateway to two peer VPN gateways (click to enlarge)

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 a non-GCP 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.

- In order 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.
  (https://cloud.google.com/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.
  (https://cloud.google.com/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

Note: All Cloud VPN gateways created before the introduction of HA VPN are considered Classic VPN gateways. To move from Classic VPN to HA VPN, see the detailed instructions. (https://cloud.google.com/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) (https://cloud.google.com/vpn/docs/concepts/overview#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. (https://cloud.google.com/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.

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

HA VPN	Classic VPN

	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

What's next

Creating HA VPNs

- Creating an HA VPN to a peer VPN gateway
  (https://cloud.google.com/vpn/docs/how-to/creating-ha-vpn)

- Creating Google Cloud to Google Cloud HA VPN gateways
  (https://cloud.google.com/vpn/docs/how-to/creating-ha-vpn2)

Creating Classic VPNs

- Creating a Classic VPN using dynamic routing
  (https://cloud.google.com/vpn/docs/how-to/creating-vpn-dynamic-routes)

- Creating a Classic VPN using static routing
  (https://cloud.google.com/vpn/docs/how-to/creating-policy-based-vpns)

Advanced configurations

- See Advanced Configurations (https://cloud.google.com/vpn/docs/concepts/advanced) for information on high-availability, high-throughput scenarios, or multiple subnet scenarios.

Managing existing Cloud VPNs

- Maintaining Cloud VPN (https://cloud.google.com/vpn/docs/how-to/maintaining-vpns)