Cloud VPN topologies

With Cloud VPN, your on-premises hosts communicate through one or more IPsec VPN tunnels to Compute Engine Virtual Machine (VM) instances in your project's VPC networks.

This page covers recommended topologies for HA VPN (https://cloud.google.com/vpn/docs/how-to/choosing-a-vpn). For Classic VPN (https://cloud.google.com/vpn/docs/how-to/choosing-a-vpn) topologies, see the Classic VPN topologies page (https://cloud.google.com/vpn/docs/concepts/classic-topologies).

To learn about the basic concepts of Cloud VPN, see the Cloud VPN Overview (https://cloud.google.com/vpn/docs/concepts/overview).

Overview

HA VPN supports site-to-site VPN in one of the following recommended topologies or configuration scenarios. Check with the vendor of your peer (https://cloud.google.com/vpn/docs/concepts/overview#peer-definition) VPN gateway to determine the appropriate configuration scenario to use:

- **An HA VPN gateway to peer VPN devices**. All of these topologies require two VPN tunnels from the perspective of the HA VPN gateway. Check with the vendor of your peer VPN gateway to determine which topology is most appropriate.
    - An HA VPN gateway to two separate peer VPN devices where each peer device has its own public IP address
    - An HA VPN gateway to one peer VPN device that has two separate public IP addresses

    - An HA VPN gateway to one peer VPN device that has one public IP address
- **An HA VPN gateway to an AWS virtual private gateway**, which is a peer gateway configuration with four interfaces.
- **Two HA VPN gateways connected to each other**.

Note: All peer gateway scenarios are represented in Google Cloud by a single external peer VPN resource.

Configurations that support 99.99% availability

To guarantee a 99.99% availability SLA for HA VPN connections (https://cloud.google.com/vpn/docs/concepts/overview#connection), *you must properly configure 2 or 4 tunnels* from your HA VPN gateway to your peer VPN gateway or to another HA VPN gateway.

Proper configuration means that VPN tunnels must supply adequate redundancy by connecting to all interfaces of the HA VPN gateway and to all interfaces of the peer VPN gateway or other HA VPN gateway.

Note: Receiving an end-to-end 99.99% availability SLA also depends on proper configuration of the peer VPN gateway.

Each of the following sections covers how to configure tunnels on both ends of the VPN connection to guarantee 99.99% availability.

Configuring HA VPN for more bandwidth

The preferred method of increasing bandwidth for HA VPN is scaling by doing the following:

- Adding more HA VPN gateways
- Adding more sets of active/passive tunnels (https://cloud.google.com/vpn/docs/concepts/overview/#active) between each HA VPN gateway interface and matching interfaces on a peer gateway (#2-peers) or on another HA VPN gateway (#2-gcp-gateways). You must match interfaces to get a 99.99% uptime SLA.

Scale gateways instead of deploying multiple tunnels connected to each interface of an existing HA VPN gateway (a bow-tie configuration),

You can connect multiple HA VPN gateways to the same peer VPN gateway (external VPN gateway resource) with as many additional tunnels as the quotas and limits for Cloud VPN (https://cloud.google.com/vpn/quotas) allow.

Below is an example of an HA VPN gateway with 10 Gbps throughput, using the following Google Cloud resources:

- 1 Cloud Router
- 4 HA VPN gateways with two tunnels apiece, for a total of 8 VPN tunnels
- 8 total BGP sessions

This configuration assumes an active/passive MED configuration for BGP sessions attached to `interface 0` and `interface 1` respectively on each gateway. That is, four `interface 0` tunnels are active and four `interface 1` tunnels are passive.

HA VPN to peer VPN gateways

There are three typical peer gateway configurations for HA VPN:

- An HA VPN gateway to two separate peer VPN devices, each with its own IP address
- An HA VPN gateway to one peer VPN device that uses two separate IP addresses
- An HA VPN gateway to one peer VPN device that uses one IP address

To set up any of these configurations, see Creating an HA VPN to a peer VPN gateway (https://cloud.google.com/vpn/docs/how-to/creating-ha-vpn).

Two peer VPN devices

If your peer-side gateway is hardware based, having a second peer-side gateway provides redundancy and failover on that side of the connection. A second physical gateway allows you to take one of the gateways offline for software upgrades or other scheduled maintenance. It also protects you in case of a failure in one of the devices.

In this topology, one HA VPN gateway connects to two peer devices. Each peer device has one interface and one public IP address. The HA VPN gateway uses two tunnels, one tunnel to each peer device.

In Google Cloud, the `REDUNDANCY_TYPE` for this configuration takes the value `TWO_IPS_REDUNDANCY`.

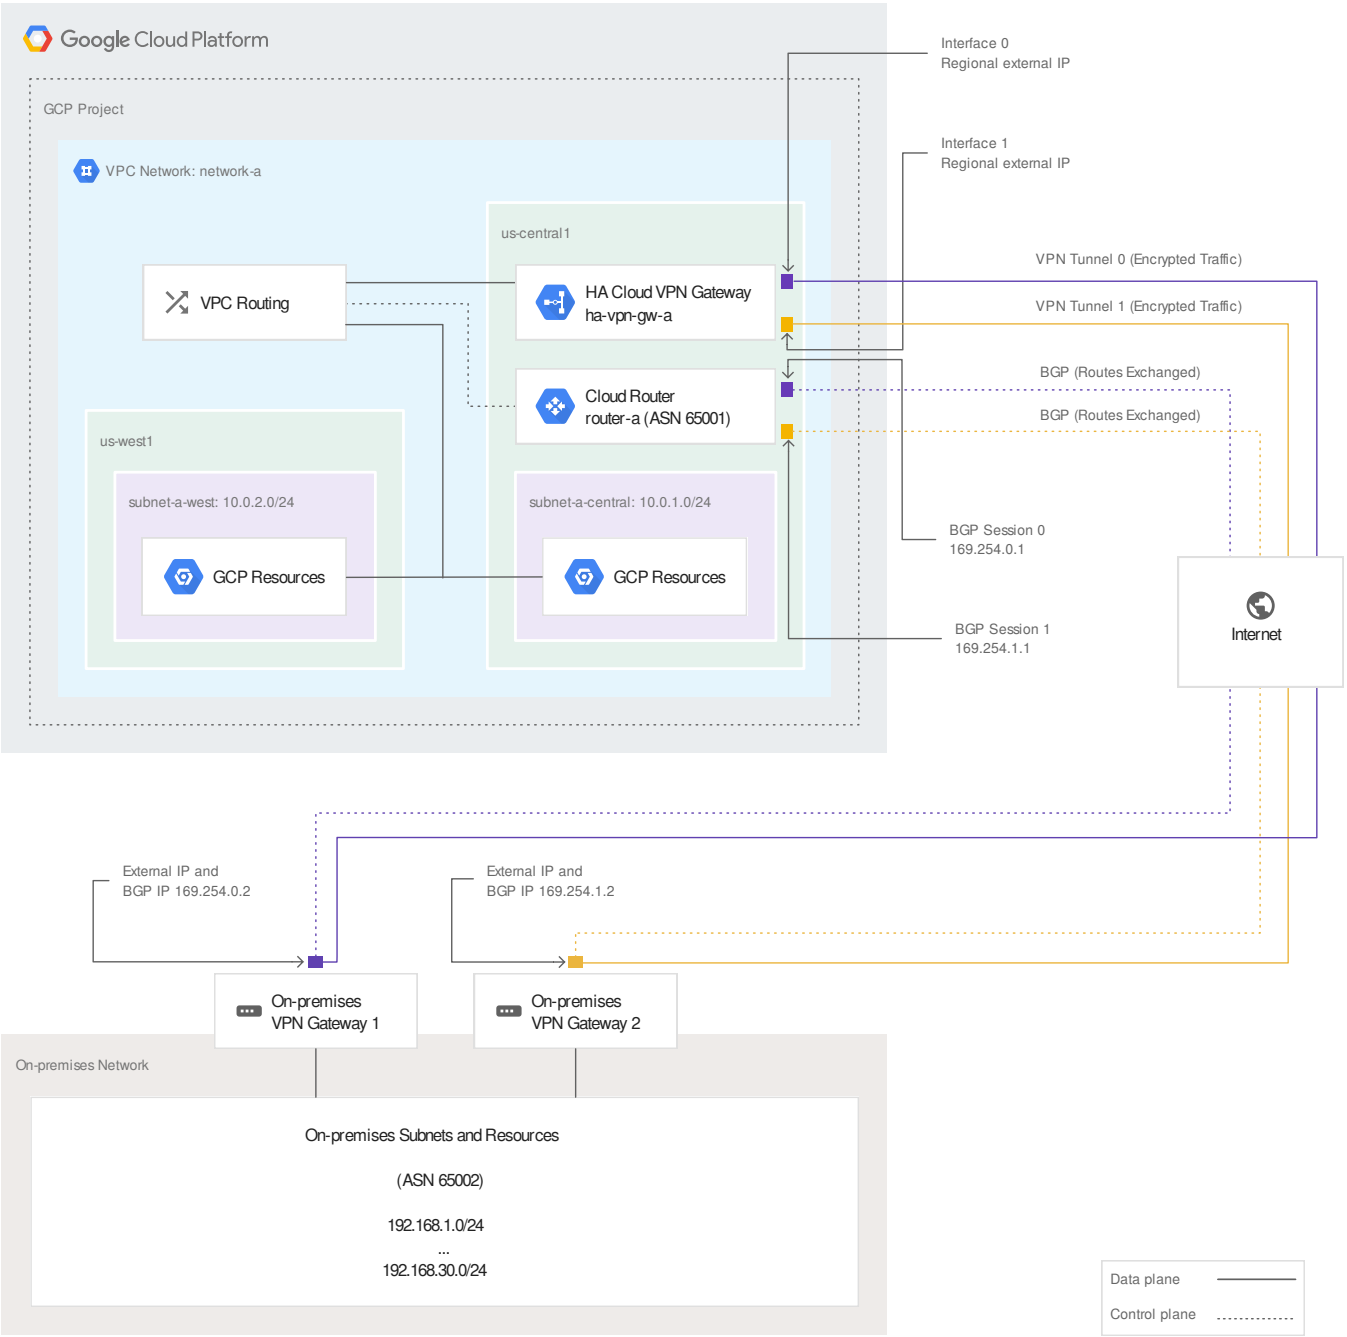

(https://cloud.google.com/vpn/images/ha-vpn-gcp-to-on-prem-2-a.svg)
HA VPN to two peer (on-premises) devices (click to enlarge)

One peer VPN device with two IP addresses

This topology describes one HA VPN gateway that connects to one peer device that has two separate public IP addresses. The HA VPN gateway uses two tunnels, one tunnel to each public IP address on the peer device.

In Google Cloud, the `REDUNDANCY_TYPE` for this configuration also takes the value `TWO_IPS_REDUNDANCY`.

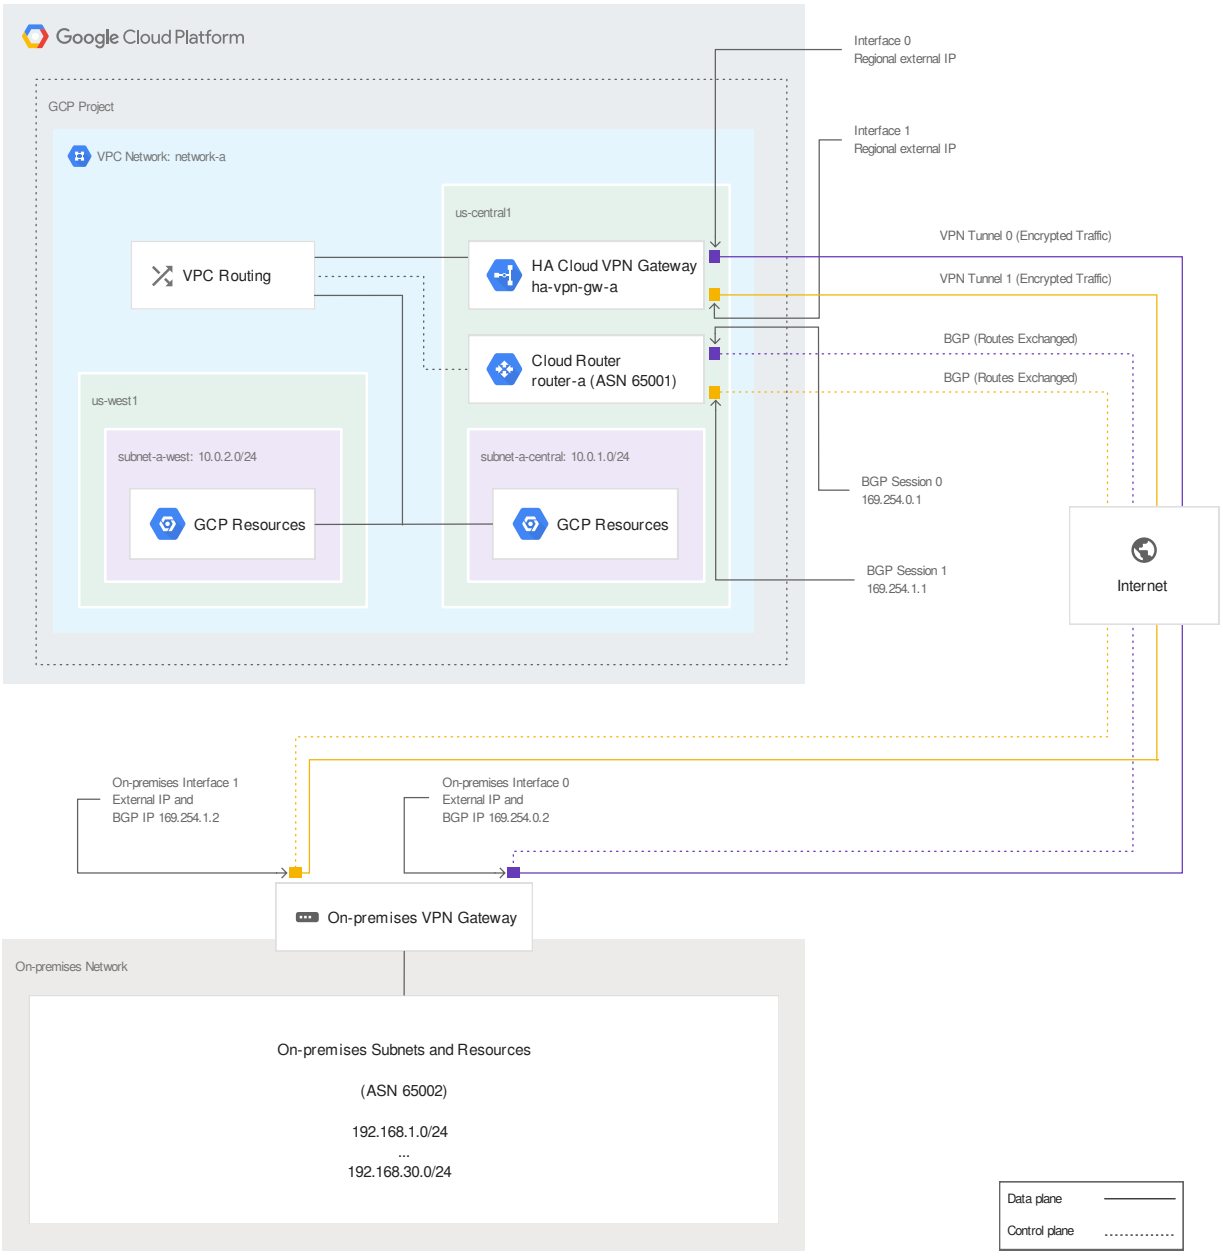

(https://cloud.google.com/vpn/images/ha-vpn-gcp-to-on-prem-2-b.svg)

HA VPN to one peer (on-premises) device with two IP addresses (click to enlarge)

One peer VPN device with one IP address

This topology describes one HA VPN gateway that connects to one peer device that has one public IP address. The HA VPN gateway uses two tunnels, both tunnels to the single public IP address on the peer device.

In Google Cloud, the `REDUNDANCY_TYPE` for this configuration takes the value `SINGLE_IP_INTERNALLY_REDUNDANT`.

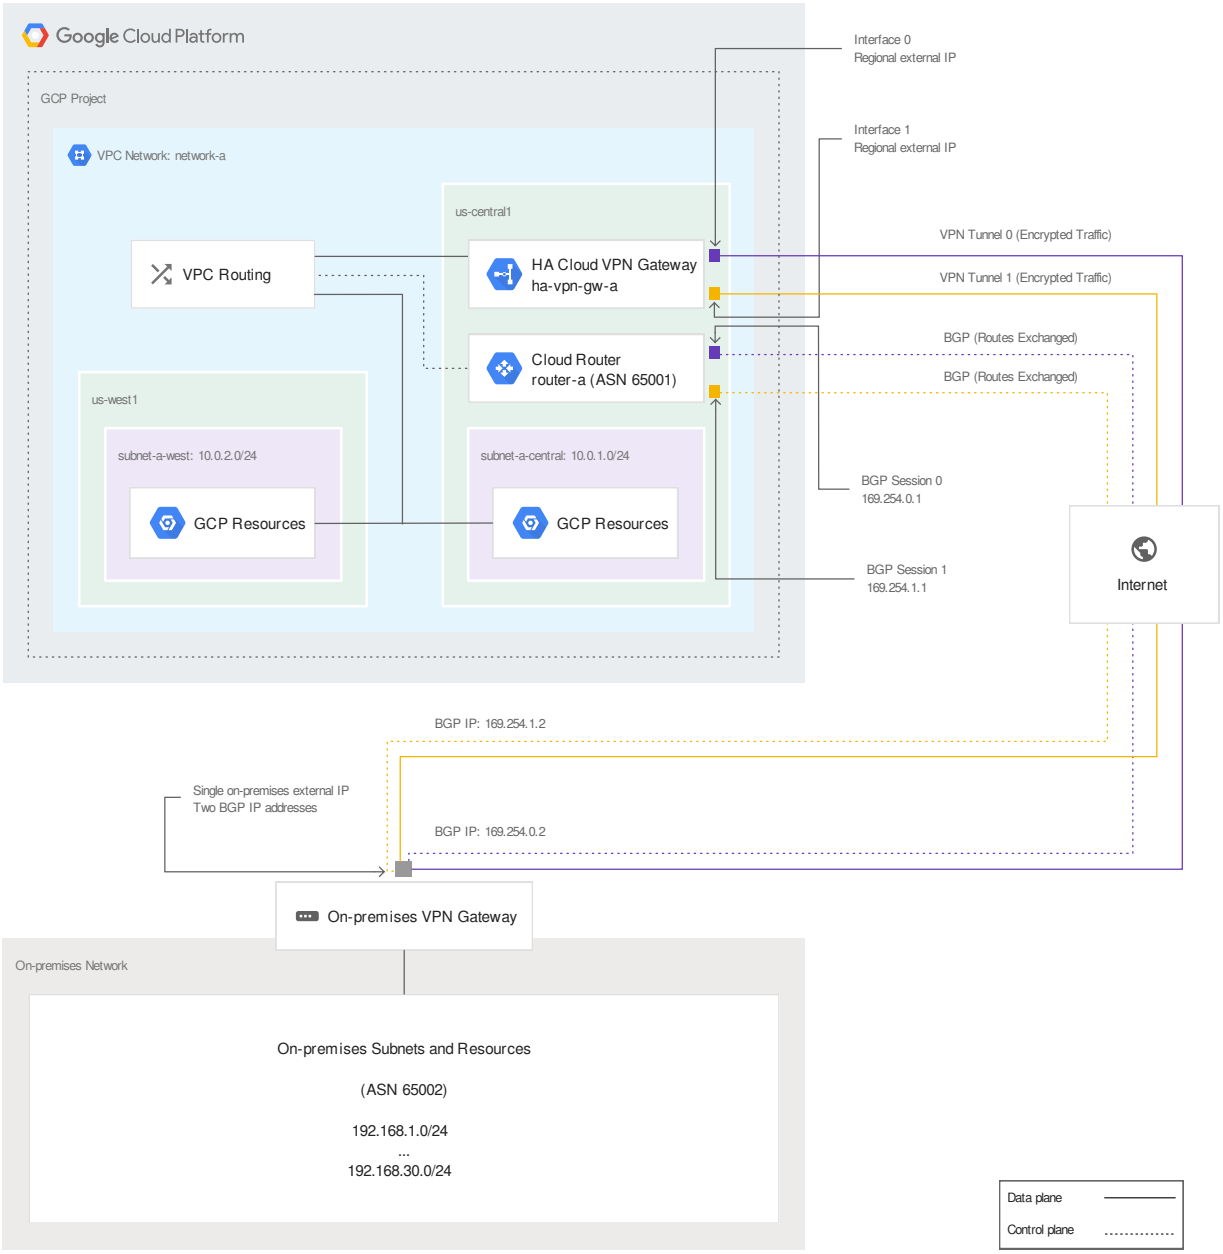

(https://cloud.google.com/vpn/images/ha-vpn-gcp-to-on-prem-2-c.svg)

HA VPN to one peer (on-premises) device with one IP address (click to enlarge)

AWS peer gateways

When configuring an HA VPN external VPN gateway to Amazon Web Services (AWS), the supported topology requires two AWS Virtual Private Gateways, `A` and `B`, each with two public IP addresses. This topology yields four public IP addresses total in AWS: `A1`, `A2`, `B1`, and `B2`.

Known issue: When configuring VPN tunnels to AWS, use the IKEv2 encryption protocol and select fewer transform sets on the AWS side, otherwise the Cloud VPN tunnel can fail to rekey. For example, select a combination of single Phase 1 and Phase 2 encryption algorithms, integrity algorithms, and DH group numbers.

This rekeying issue is caused by a large SA payload size for the default set of AWS transform sets. This large payload size results in IP fragmentation of IKE packets on the AWS side, which Cloud VPN does not support.

1. Configure the four AWS IP addresses as a single external HA VPN gateway with `FOUR_IPS_REDUNDANCY`, where:
    - AWS IP `0=A1`
    - AWS IP `1=A2`
    - AWS IP `2=B1`
    - AWS IP `3=B2`
2. Create four tunnels on the HA VPN gateway to meet the 99.99% SLA. using the following configuration:
    - HA VPN interface 0 to AWS interface 0
    - HA VPN interface 0 to AWS interface 1
    - HA VPN interface 1 to AWS interface 2
    - HA VPN interface 1 to AWS interface 3

Overview of high-level configurations steps to set up HA VPN with Amazon Web Services (AWS):

1. Create the HA VPN gateway and a Cloud Router. This creates 2 public IP addresses on the GCP side.

2. Create two AWS Virtual Private Gateways. This creates 4 public addresses on the AWS side.
3. Create two AWS Site-to-Site VPN connections and customer gateways, one for each AWS Virtual Private Gateway. Specify a non-overlapping link-local Tunnel IP Range for each tunnel, 4 total. For example, 169.254.1.4/30.
4. Download the AWS configuration files for the generic device type.
5. Create four VPN tunnels on the HA VPN gateway.
6. Configure BGP sessions on the Cloud Router using the BGP IP addresses from the downloaded AWS configuration files.

Guaranteeing 99.99% availability

To meet the 99.99% SLA on the GCP side, there must be a tunnel from each of the two interfaces on the HA VPN gateway to the corresponding interface(s) on the peer gateway.

If the peer gateway has two interfaces, then configuring two tunnels, one from each peer interface to each HA VPN gateway interface, meets the requirements for the 99.99% SLA. A full mesh configuration, is *not* required for 99.99% SLA *on the GCP side*. In this case, a full mesh is defined as two tunnels from each HA VPN interface to each of the two interfaces on the peer gateway, for a total of four tunnels from the Google Cloud side. See the documentation for your peer (on-premises) VPN device, or contact your VPN vendor, to confirm if they recommend a full mesh configuration.

The following example provides 99.99% availability:

In this configuration, tunnels on each of the following interfaces on the HA VPN gateway match the corresponding interfaces on the peer gateway or gateways:

- HA VPN `interface 0` to peer `interface 0`
- HA VPN `interface 1` to peer `interface 1`

Examples are shown in the drawings for two peer devices, two interfaces (#2-device-2-peers-drawing) and one peer device, two interfaces (#1-device-2-peers-drawing).

If there is only one peer interface on one peer gateway, each tunnel from each HA VPN gateway interface must connect to the single peer interface. An example of this is shown in the drawing for one peer device, one interface (#1-peer-drawing).

Caution: You must configure at least one tunnel on each HA VPN gateway interface to receive the 99.99% availability SLA. Configuring only one tunnel from a single HA VPN interface to a single interface on the peer gateway doesn't provide enough redundancy to meet the SLA because there is an unused interface on the HA VPN gateway, which does not have a tunnel configured on it.

The following example *doesn't* provide 99.99% availability:

- HA VPN `interface 0` to peer `interface 0`

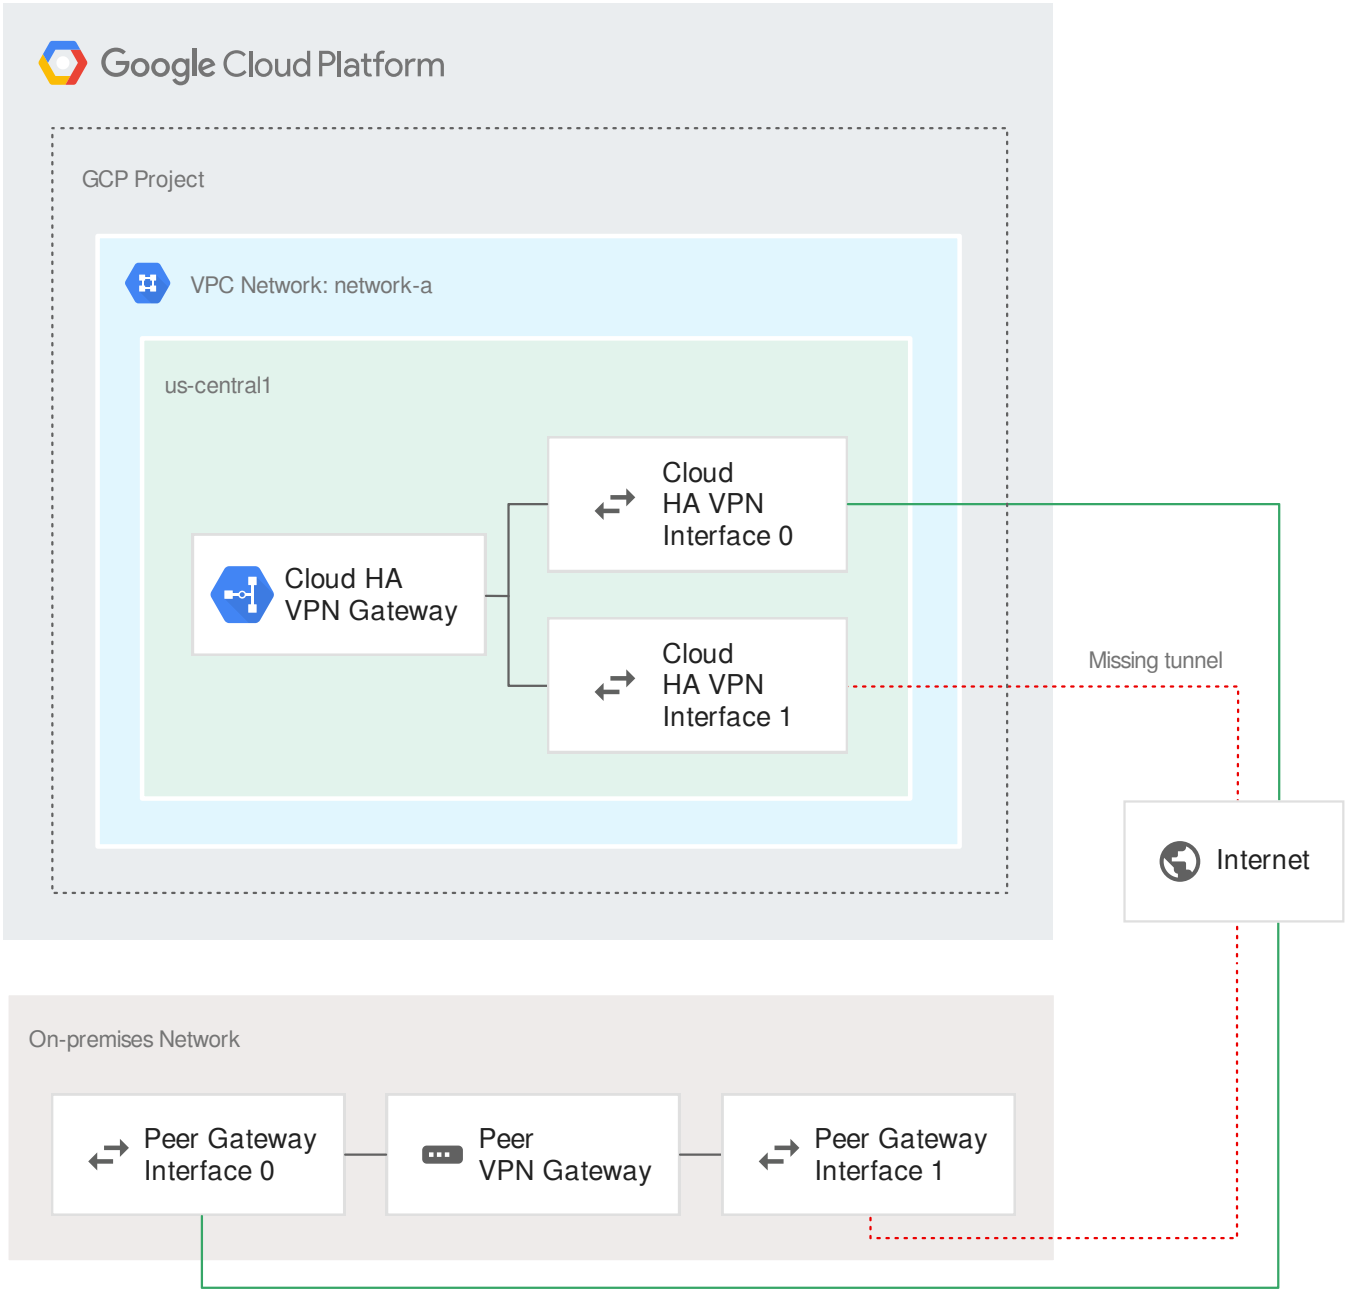

(https://cloud.google.com/vpn/images/ha-gcp-2-peer-1-tunnel.svg)

A topology that doesn't provide high availability (click to enlarge)

Google Cloud-to-Google Cloud HA VPN gateways

You can connect two Google Cloud VPC networks together using an HA VPN gateway in each network.

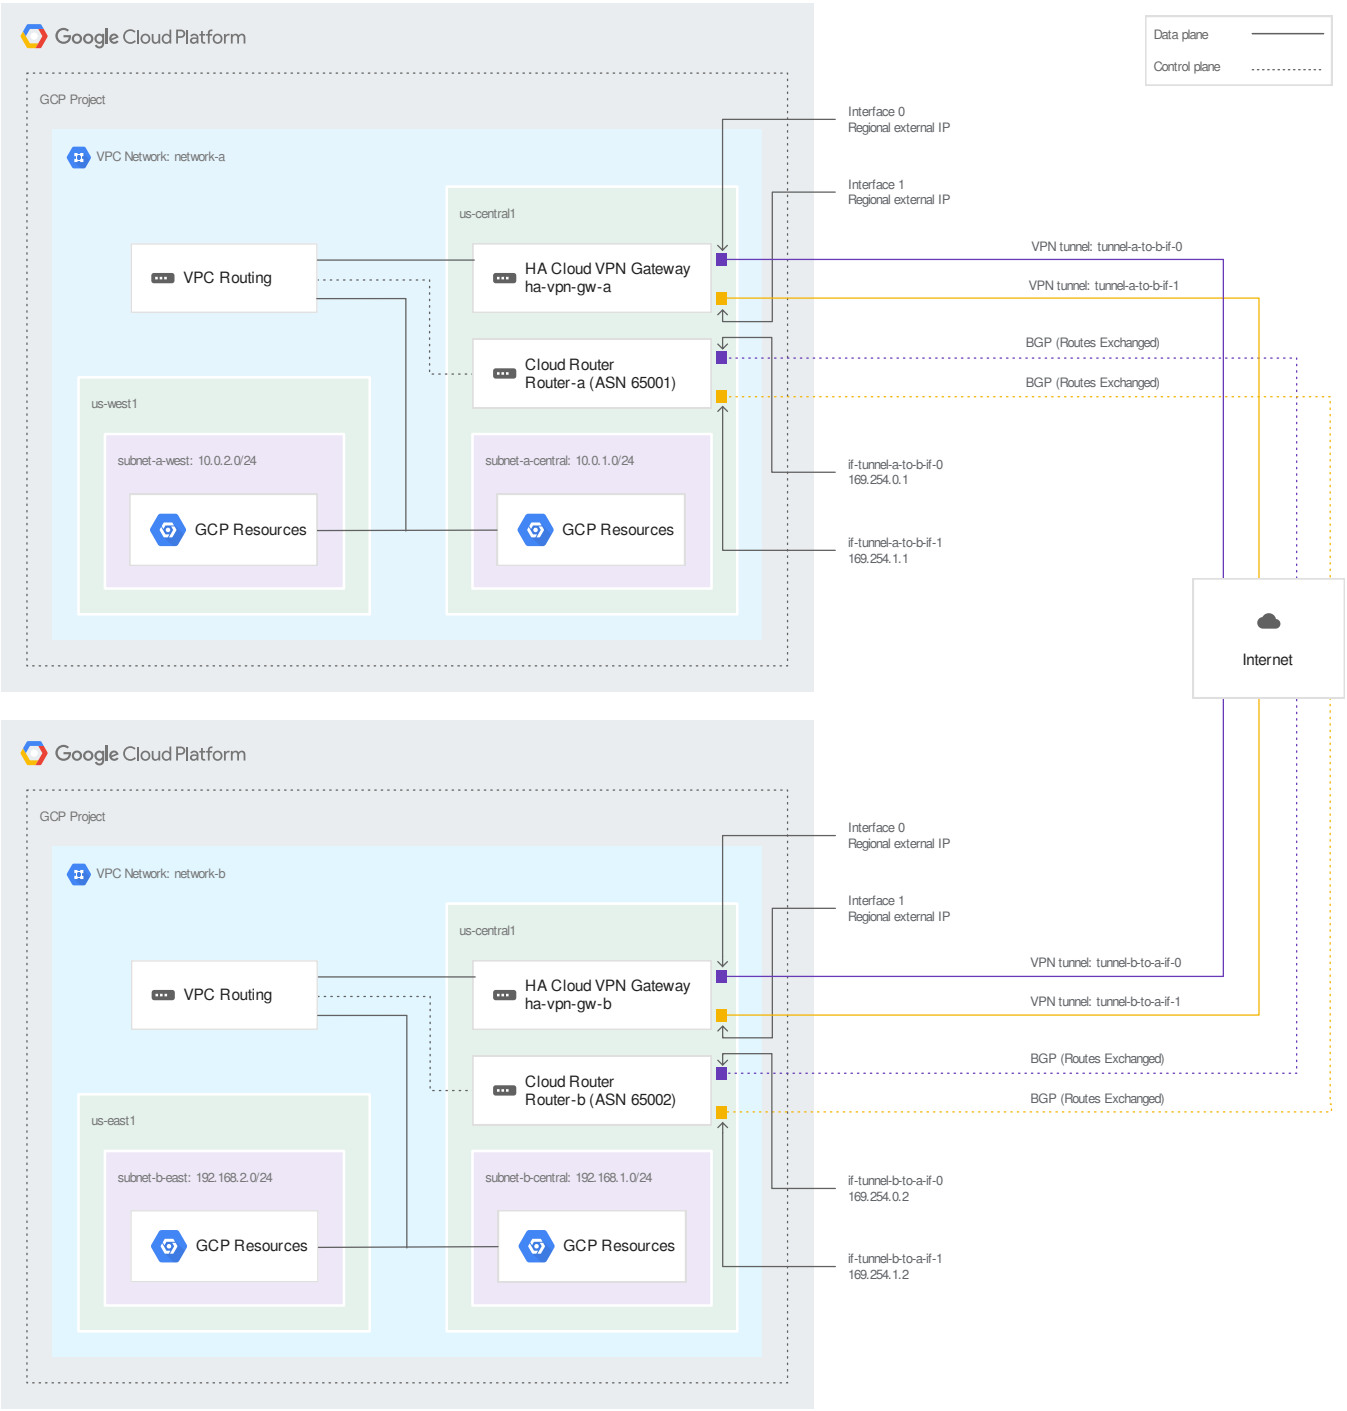

(https://cloud.google.com/vpn/images/ha-vpn-gcp-to-gcp.svg)

Google Cloud to Google Cloud HA VPN gateways (click to enlarge)

From the perspective of each HA VPN gateway, you create two tunnels so that both of the following are true:

- `interface 0` on one HA VPN gateway to `interface 0` on the other HA VPN
- `interface 1` on one HA VPN gateway to `interface 1` of the other HA VPN.

To set up this configuration, see Creating HA VPN to HA VPN gateways (https://cloud.google.com/vpn/docs/how-to/creating-ha-vpn2).

Guaranteeing 99.99% availability

Caution: Configuring only one tunnel on one interface for each HA VPN gateway does not provide a 99.99% SLA.

To provide 99.99% availability for HA VPN to HA VPN gateways , the following interfaces on both gateways must match:

- HA VPN `interface 0` to HA VPN `interface 0` and
- HA VPN `interface 1` to HA VPN `interface 1`

What's next

- Learn about the basic concepts of Cloud VPN (https://cloud.google.com/vpn/docs/concepts/overview)
- Create a custom Virtual Private Cloud network (https://cloud.google.com/vpc/docs/using-vpc#create-custom-network)
- Set up different types of Cloud VPN (https://cloud.google.com/vpn/docs/how-to/choosing-a-vpn)
- Maintain VPN tunnels and gateways (https://cloud.google.com/vpn/docs/how-to/maintaining-vpns)
- See Advanced Configurations (https://cloud.google.com/vpn/docs/concepts/advanced) for information on high-availability, high-throughput scenarios, or multiple subnet scenarios.
- View logs and monitoring metrics (https://cloud.google.com/vpn/docs/how-to/viewing-logs-metrics)
- Get troubleshooting help (https://cloud.google.com/vpn/docs/support/troubleshooting)
- Calculate network throughput (https://cloud.google.com/community/tutorials/network-throughput).

Previous

← Cloud VPN and other hybrid connectivity solutions
(https://cloud.google.com/vpn/docs/concepts/choosing-a-hybrid-solution)

Next
Networks and tunnel routing →
(https://cloud.google.com/vpn/docs/concepts/choosing-networks-routing)