Google Cloud

Cloud WAN Solution Guide

October 2025

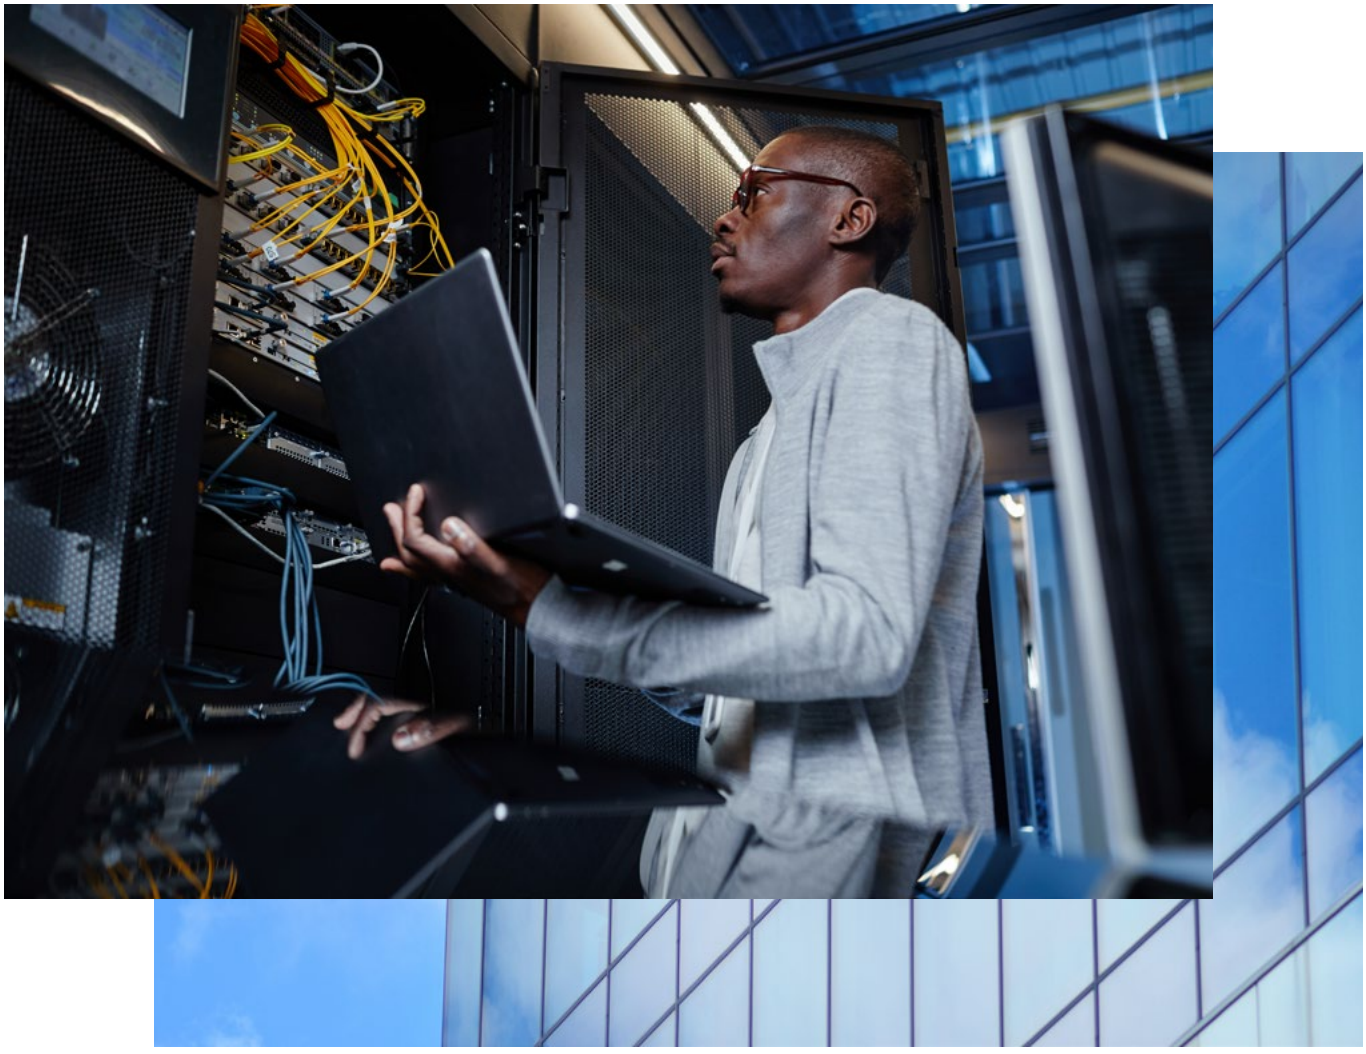

Table of contents

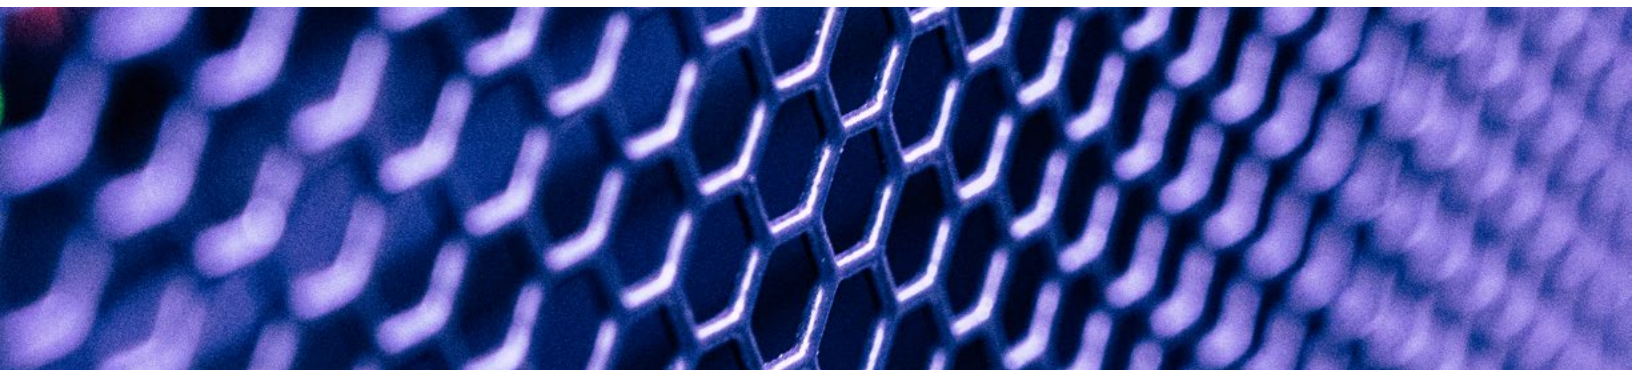

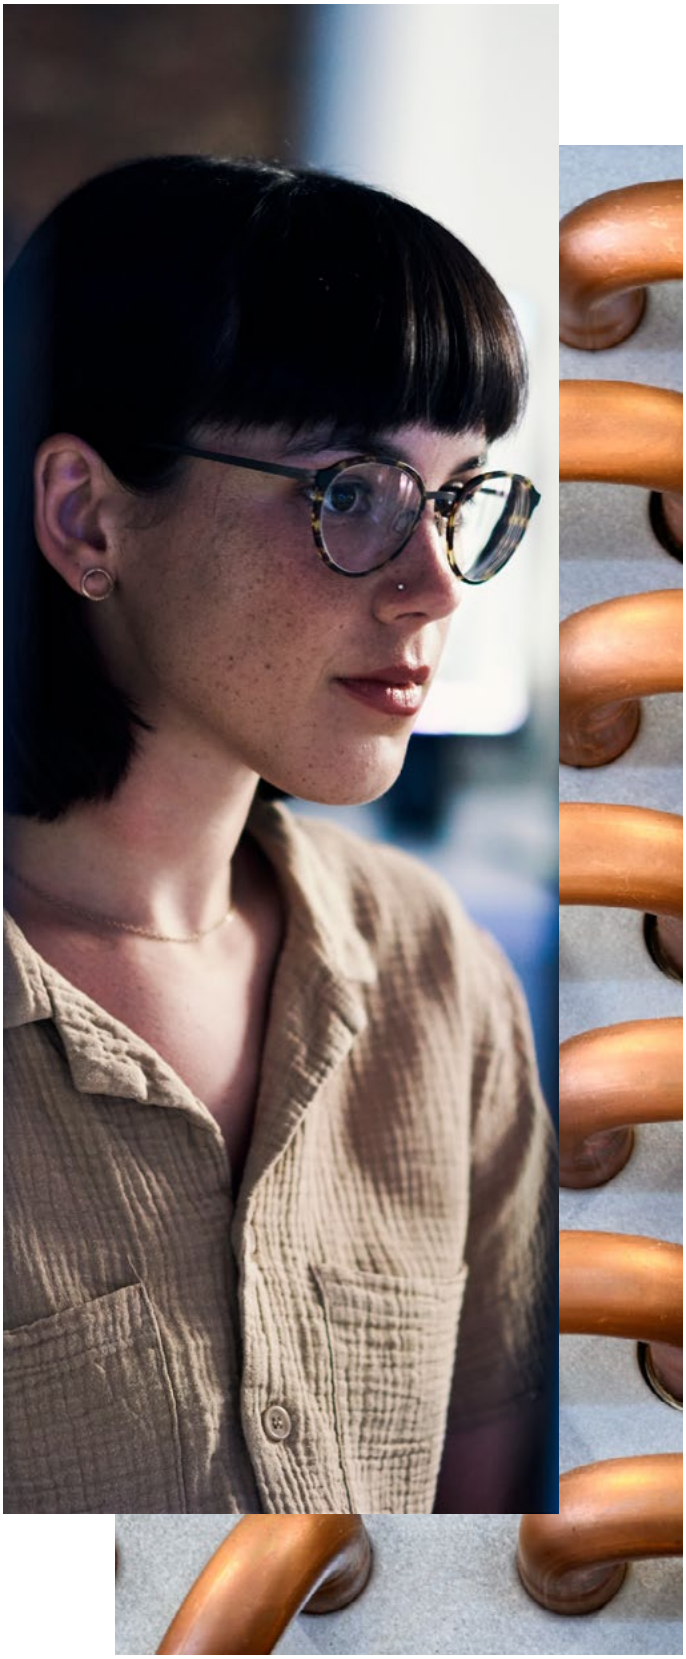

01

Executive summary

This document provides a comprehensive guide for cloud practitioners on implementing recommended architectures for Cloud WAN solutions. It details how to leverage a combination of Google Cloud networking products, including Cloud Interconnect, Cross-Cloud Interconnect, Cross-Site Interconnect, Network Connectivity Center (NCC), and integrations with third-party network virtual appliance (NVA) vendors, to build robust and scalable enterprise WAN architectures.

02

Overview of Cloud WAN

Cloud WAN is a fully managed, reliable, and secure enterprise backbone designed to modernize and transform wide area network (WAN) architectures. It leverages Google's planet-scale network, optimized for application performance, positioning itself as the third core pillar of Google's Cross-Cloud Network platform.

The Cross-Cloud Network provides the foundational infrastructure for interconnecting diverse networks in a global multicloud environment and for the optimal deployment of security services. Key capabilities enabled by the Cross-Cloud Network include:

01

Building distributed applications

Facilitating the efficient composition of applications by leveraging best-in-class services hosted across cloud providers and on-premise environments.

02

Delivering a Global Front End

Securely and efficiently serving web origins distributed across hybrid multicloud infrastructures.

03

Connecting Cloud WAN securely

Enabling secure connectivity for a hybrid workforce by utilizing the global scale, reliability, and ubiquity of the Google network.

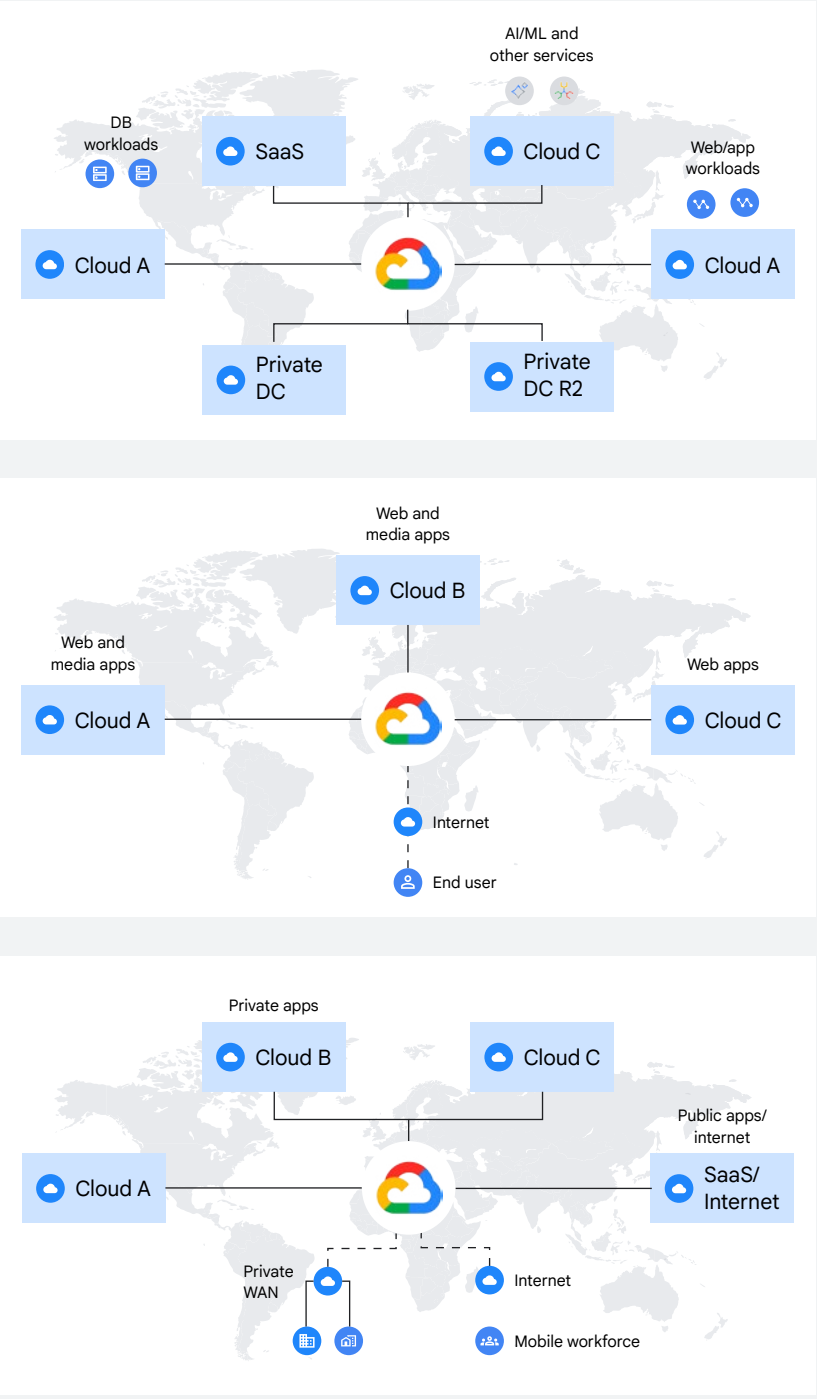

Figure 2.1: Pillars of Cross-Cloud Network

Cloud WAN primarily supports two critical use cases: establishing high-performance connectivity between geographically dispersed data centers and seamlessly connecting branch and campus environments over Google Cloud's Premium Tier network. The subsequent sections elaborate on the specific Cloud WAN solutions for each of these use cases.

In the next two sections, we elaborate on each of the use cases unblocked using the different Cloud WAN solutions.

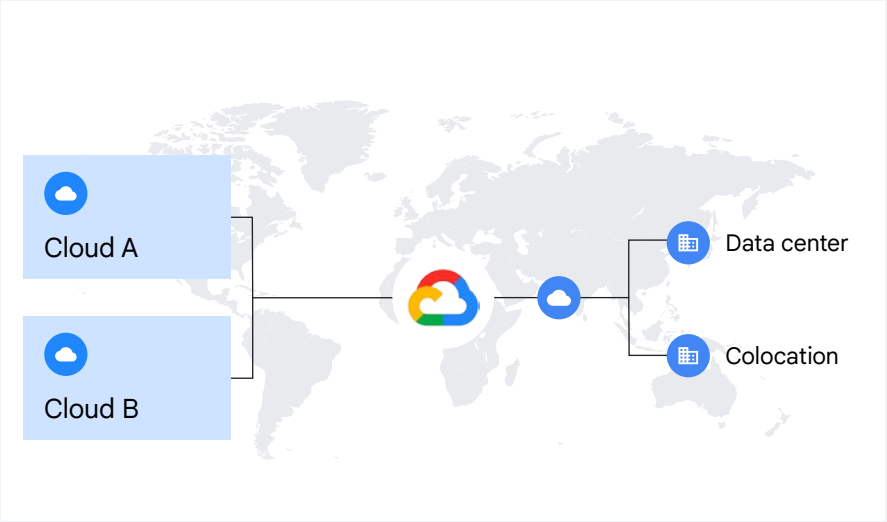

Figure 2.2: Cloud WAN use case 1

Enable high-performance connectivity

Use cases: Data center interconnect, global WAN across sites

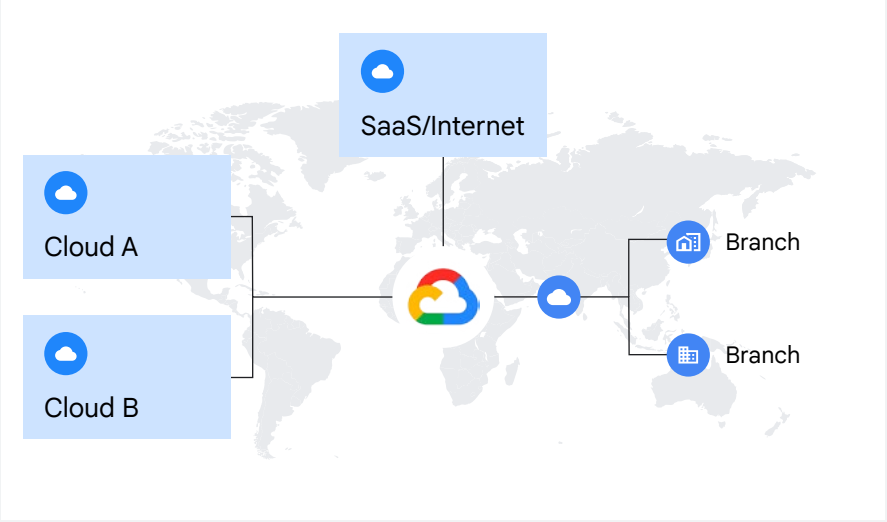

Figure 2.3: Cloud WAN use case 2

Migrate branch and campus networks

Use cases: Branch and campus cloud onramp

03

Enable high-performance connectivity

A primary driver for enabling high-performance connectivity with Cloud WAN is to allow global customers to reliably and efficiently transfer large volumes of data across their data centers and/or other cloud providers. Customers typically adopt one of the following architectural approaches:

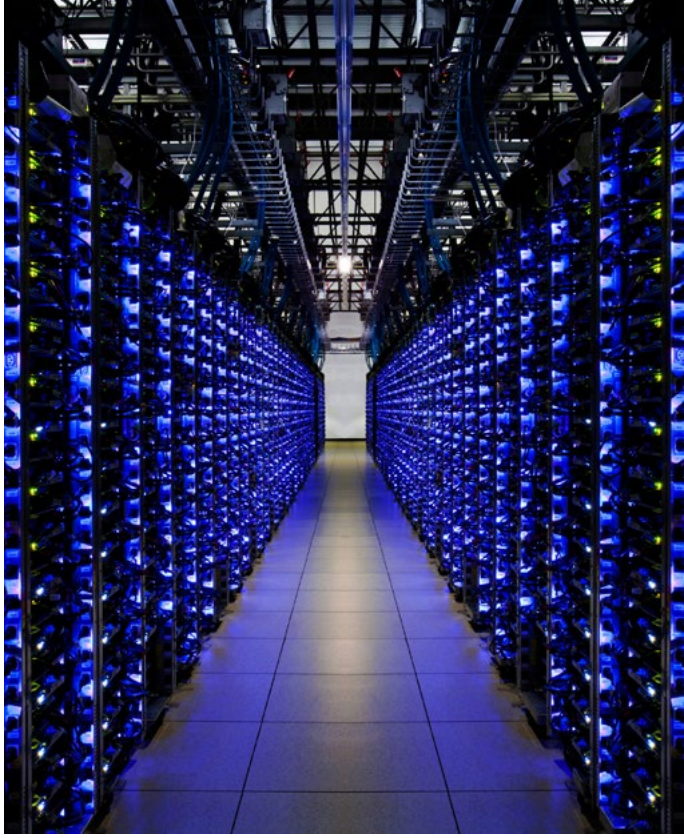

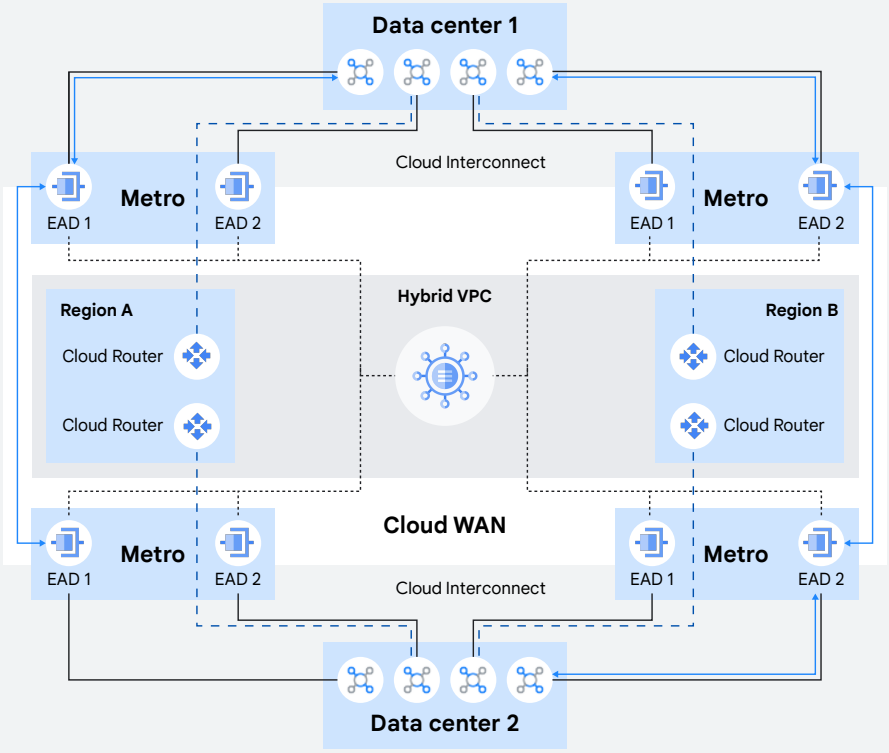

Figure 3.1: Site-to-site connectivity for data centers using NCC

a. Site-to-site connectivity for data centers using NCC

This design leverages Cloud Interconnect in conjunction with Network Connectivity Center (NCC) to facilitate high-bandwidth, site-to-site data transfer between customer sites.

Cloud Interconnect integration: Dedicated Interconnects provide high-bandwidth, private connectivity from the customer's data center directly to Google Cloud.

High availability design: To achieve a 99.99% availability service level agreement (SLA), the recommended architecture as illustrated in Figure 3.1 requires a minimum of four Dedicated Interconnect connections. These connections must be geographically distributed across two metropolitan areas, with two connections established in each metro. This design mitigates risks associated with single points of failure at both the physical interconnect and metropolitan levels.

BGP peering configuration: Customers are required to establish Border Gateway Protocol (BGP) peering between their on-premises data center router(s) and the respective Cloud Router instances in each Google Cloud region. This BGP session facilitates the exchange of route prefixes over the VLAN attachments configured on Cloud Interconnect.

NCC hub integration: The VLAN attachments associated with each Cloud Interconnect are attached as spokes to the NCC hub.

Site-to-site data transfer enablement: To enable seamless traffic routing between the connected data centers, site-to-site data transfer functionality must be explicitly enabled on the NCC hub. This configuration allows traffic to traverse securely and efficiently between customer data centers via Cloud Interconnect, leveraging the optimized Google backbone network for inter-site communication.

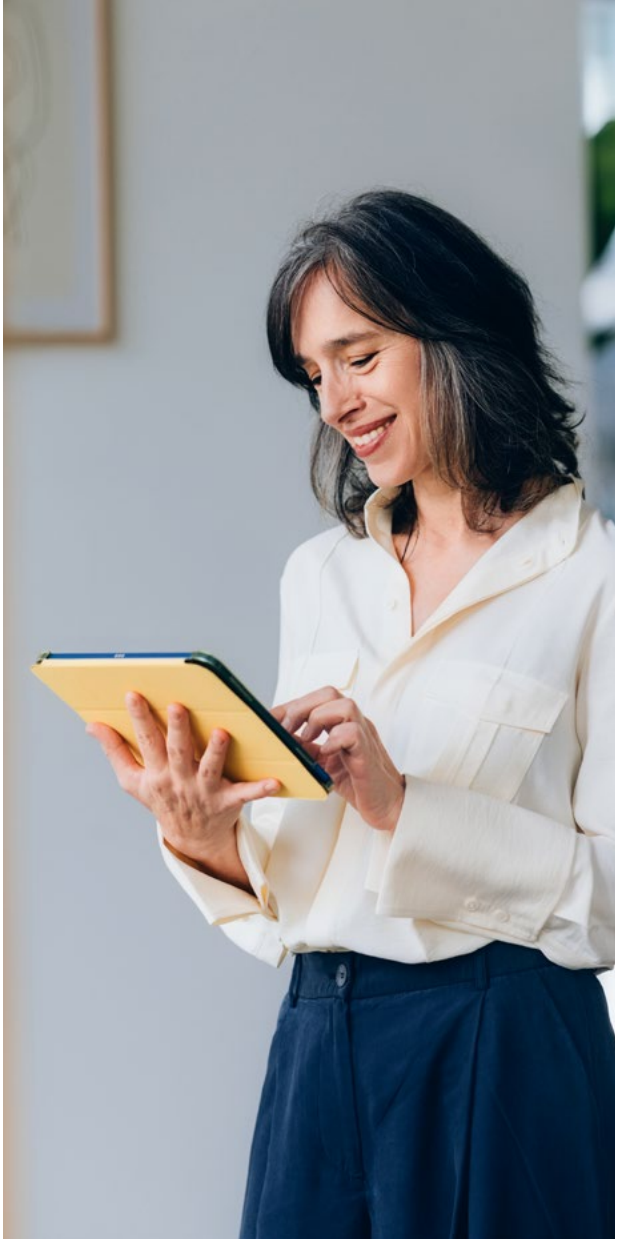

Key considerations:

For Cloud Routers in different regions to exchange routes and enable full global reachability, global dynamic routing mode must be enabled in your Virtual Private Cloud (VPC) network.

Site-to-site data transfer is only available in supported locations. Refer to the official Google Cloud documentation for Network Connectivity Center supported locations for more information.

b. Site-to-site connectivity for data centers using Cross-Site Interconnect

This design offers an alternative to NCC-based connectivity, leveraging Cross-Site Interconnect (CSI) for high-bandwidth, site-to-site data transfer directly between customer sites.

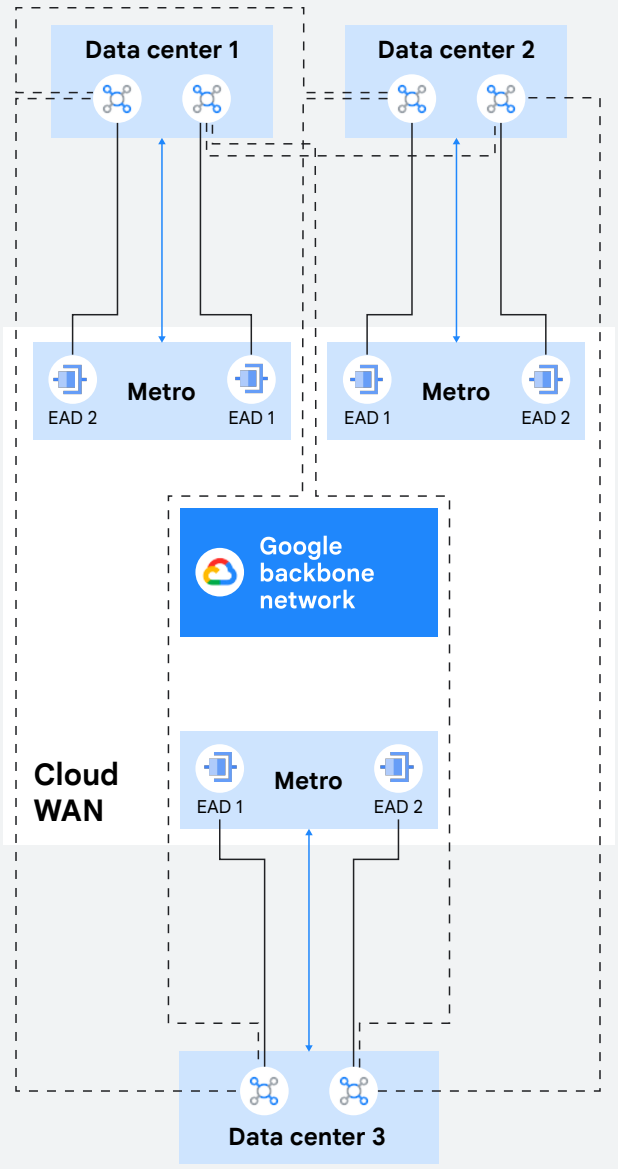

Figure 3.2: Site-to-site connectivity for data centers using CSI

Direct private connectivity: CSI provides private, high-bandwidth Layer 2 connectivity directly between customer data centers utilizing the Google backbone network.

Simplified deployment: A key differentiator of Cross-Site Interconnect from Cloud Interconnect is its transparent Layer 2 nature, which eliminates the requirement for BGP peering with Cloud Routers or the deployment of Network Connectivity Center. This simplifies network configuration and management.

Pseudowire deployment modes: Customers can deploy individual pseudowires from their on-premises routers over the Cross-Site Interconnect using one of two modes:

- Port mode: Transparently forwards all traffic on a physical port to a single destination, irrespective of VLAN tags or other headers. This provides a simple, direct link.
- VLAN mode: Enables the creation of multiple virtual networks (VLANs) on a single physical port, allowing for connections to distinct sites or network segments.

Full mesh connectivity with ring topology: As illustrated in Figure 3.2, this design can configure three distinct wire groups in VLAN mode to achieve full-mesh connectivity across three data centers by employing a ring topology. This provides robust inter-site communication.

High availability: High availability site-to-site connectivity is inherently ensured through the deployment of redundant Cross-Site Interconnect connections.

c. Cloud-to-cloud connectivity using Cross-Cloud Interconnect

This advanced design leverages a combination of Cloud Interconnect and Cross-Cloud Interconnect with NCC to establish comprehensive site-to-site data transfer, encompassing both on-premises data centers and connections to other cloud providers.

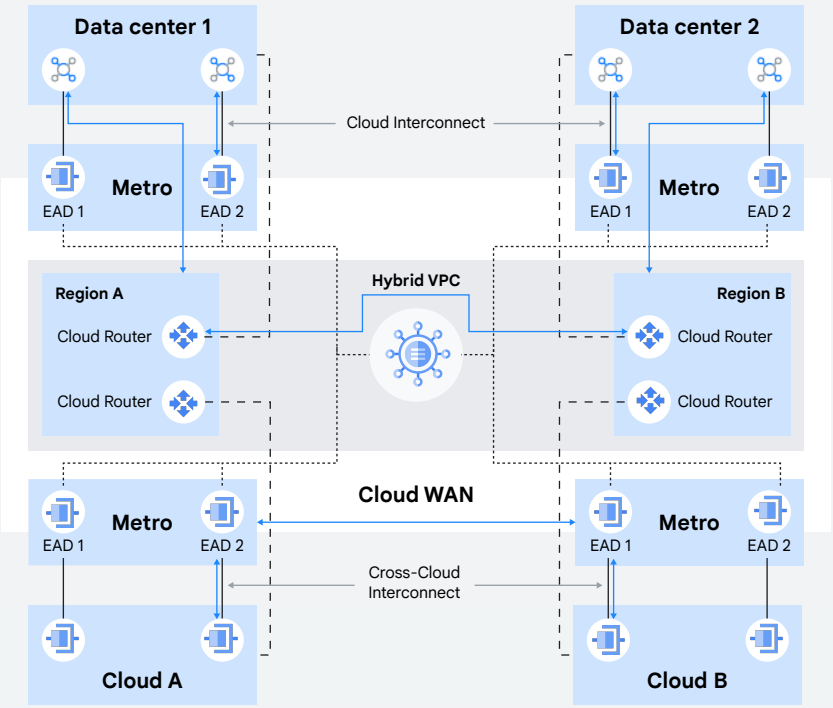

Figure 3.3: Cloud-to-cloud connectivity using Cross-Cloud Interconnect

Cross-Cloud Interconnect (CCI) functionality: Cross-Cloud Interconnect provides high-bandwidth, private connectivity specifically between Google Cloud and other hyperscale cloud providers.

Managed physical connectivity: Google Cloud simplifies deployment by provisioning the physical connectivity between its network and the other cloud provider's network in supported colocation facilities. This eliminates the need for the customer to procure and deploy their own physical hardware in these facilities.

VLAN attachments and BGP peering: Customers initiate requests for physical connections in the desired colocation facility(s) and subsequently configure corresponding VLAN attachments over these physical connections for the respective cloud provider(s). These VLAN attachments are crucial for establishing BGP peering between Cloud Router and the equivalent BGP peer in the other cloud provider's environment.

NCC integration for centralized control: The configured VLAN attachments are then attached as spokes to the central Network Connectivity Center (NCC) hub.

Inter-cloud data transfer: To enable seamless traffic routing and exchange between the connected cloud providers, site-to-site data transfer must be activated on the NCC hub. This configuration ensures that traffic can traverse efficiently and securely between the different cloud providers by leveraging the high-performance Google backbone network.

Availability supported: The design illustrated in Figure 3.3 supports a 99.9% availability SLA for site-to-site data transfer between customer data center(s) and/or cloud provider(s) when utilizing an NCC hub with site-to-site data transfer enabled.

04

Migrate branch and campus network

This use case focuses on enabling customers to migrate their branch and campus networks by leveraging Google Cloud's Premium Tier network. This effectively serves as a robust and high-performance on-ramp for these remote sites to access private and public cloud resources, as well as SaaS applications.

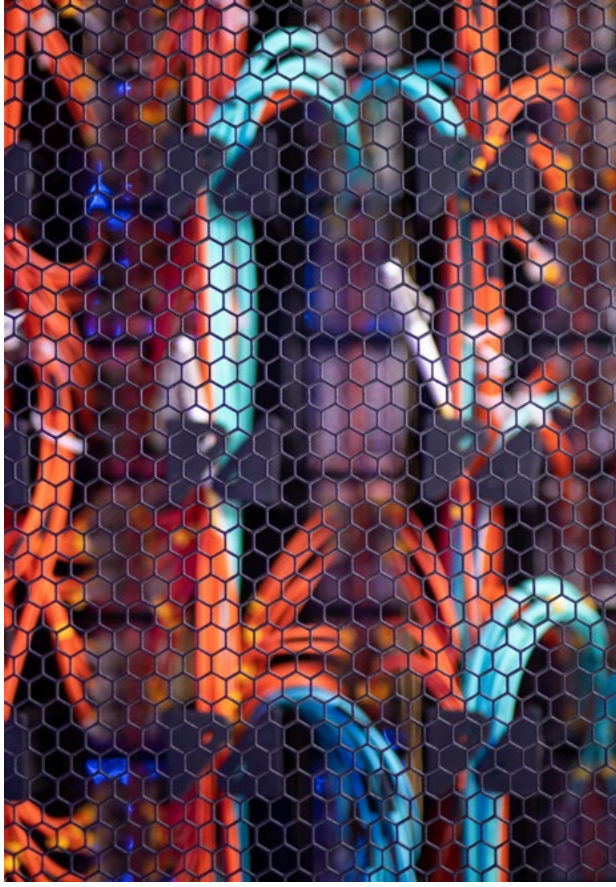

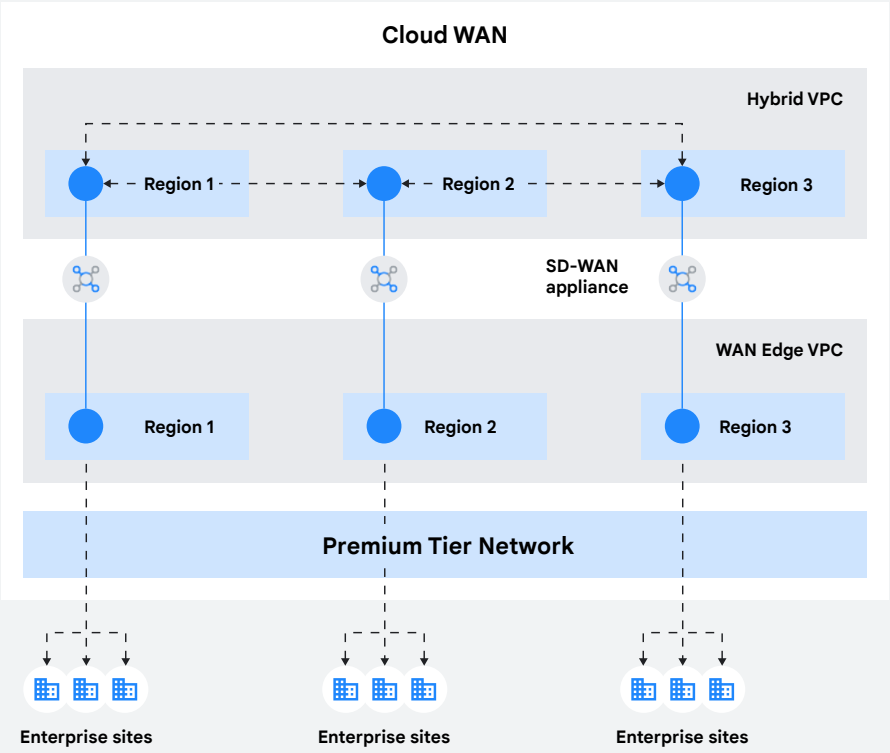

Figure 4.1: Site-to-site connectivity for branch/campus networks

a. Site-to-site connectivity for branch/campus networks

This design pattern utilizes Google Cloud's Premium Tier network in conjunction with a multi-NIC network virtual appliance (NVA) to establish secure and performant site-to-site connectivity between geographically dispersed branch and campus networks over the internet.

Network segregation with multi-NIC NVA: The multi-NIC NVA is strategically deployed to enable clear segregation between trusted and untrusted network zones, allowing for granular routing policies.

Dual-VPC deployment model: Consequently, the NVA is typically deployed across two distinct VPCs:

- **WAN Edge Global VPC:** This VPC is dedicated to terminating untrusted or internet-based connectivity from the branch/campus sites.
- **Hybrid Global VPC:** This VPC is used for configuring the secure overlay network that extends into Google Cloud and facilitates internal routing.

Extending connectivity: Customers can extend their existing connectivity from branch/campus networks by deploying NVAs within a VPC.

Overlay network for scale and isolation: The design illustrated in Figure 4.1 uses an overlay network, established from the branch/campus networks to the NVAs and subsequently between the NVAs within Google Cloud. This overlay network is an extension of the customer's SD-WAN footprint in the cloud, established using IPsec/GRE protocol or a proprietary vendor tunneling protocol. We recommended a GRE tunnel for maximum performance. The overlay network design provides significant advantages:

- It allows customers to exceed the maximum supported dynamic route advertised limit on Cloud Routers, which is critical for large deployments with many routes.
- This design is particularly beneficial for customers requiring strict Layer 3 segregation using Virtual Routing and Forwarding (VRF) between different branch/campus networks.

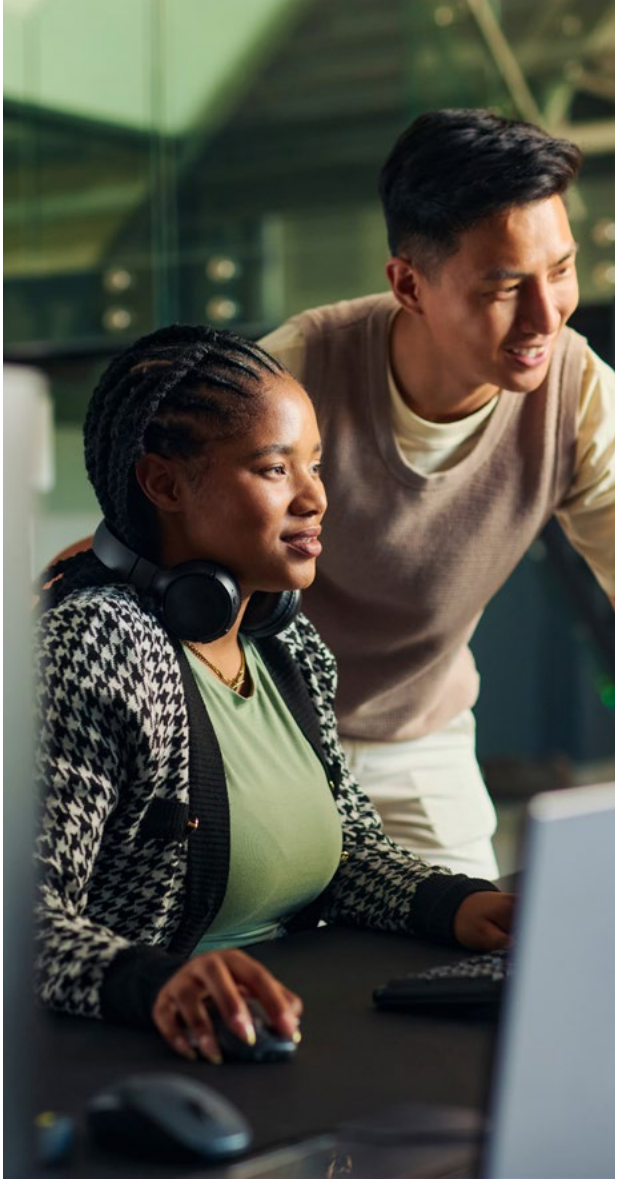

Key considerations:

- The deployment of NVA(s) to configure the overlay network may incur additional licensing and compute charges from the NVA vendor.
- Premium Tier network egress is charged at internet data transfer rates.
- Additional inter-region data transfer charges may apply for traffic traversing across different regions within the VPC network.

b. Branch/campus networks to Google Cloud connectivity

This design leverages Premium Tier network, multi-NIC network virtual appliance (NVA), and Network Connectivity Center to establish site-to-cloud connectivity between branch/campus networks over the internet to Google Cloud workload VPCs. The multi-NIC NVA enables trusted and untrusted network segregation and allows for granular routing. Hence it is deployed across 2 VPCs namely WAN Edge VPC for terminating the untrusted or over the internet connectivity and Hybrid VPC for adding as router appliance spoke of NCC hub for trusted or private connectivity to other workload VPCs.

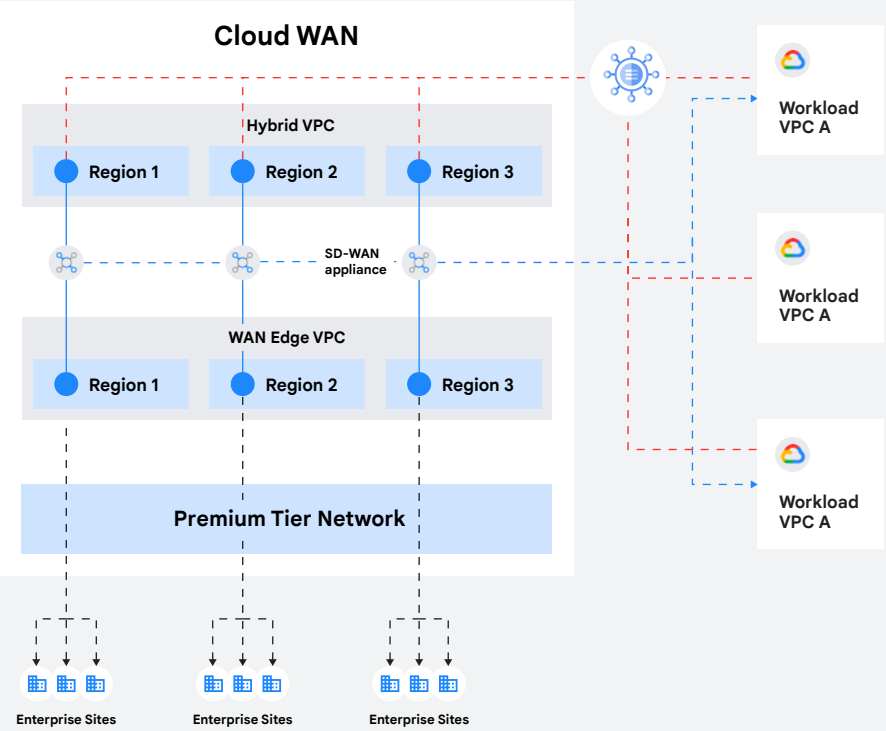

Figure 4.2(a): Branch/campus networks to Google Cloud connectivity

The design illustrated in Figure 4.2(a) extends the previous concept by leveraging Google Cloud's Premium Tier network, a multi-NIC NVA, and Network Connectivity Center to establish robust site-to-cloud connectivity. This enables branch/campus networks to securely connect over the internet to Google Cloud workload VPCs.

NVA-based network segregation and deployment: The multi-NIC NVA remains central to enabling trusted and untrusted network segregation and granular routing. It is deployed across two VPCs:

- **WAN Edge VPC:** For terminating untrusted (internet) connectivity from branch/campus sites.
- **Hybrid VPC:** This acts as a router appliance spoke of the NCC hub, providing trusted (private) connectivity to other workload VPCs.

Centralized NCC integration: Workload VPCs are added as VPC spokes to the NCC hub. The NVA(s) in the Hybrid VPC must also be added as router appliance spokes of the NCC hub, necessitating the setup of BGP peering between the NVA and the Google Cloud Router.

Advanced prefix control: This configuration offers granular control over the number of route prefixes advertised from the NVA to the NCC hub through the use of BGP filters.

- If the number of on-premises prefixes exceeds the maximum dynamic routes supported by NCC, customers can summarize and/or aggregate these prefixes before advertising them to the Cloud Router.
- Alternatively, customers can use BGP route policies to filter routes and advertise only a summarized route to the Cloud Router.
- If the number of prefixes remains within NCC's dynamic route limits, customers can advertise the on-premises route prefixes directly without aggregation or summarization.

Dynamic route exchange for cloud connectivity: NCC's support for dynamic route exchange allows seamless traffic routing between the NVAs and the workload VPCs using the underlying VPC network, thereby facilitating efficient site-to-cloud connectivity.

Key considerations:

Deployment of NVA(s) for overlay network configuration may entail additional licensing and compute charges.

Premium Tier network egress is billed at internet data transfer rates.

- NCC charges for spoke attachments, advanced data networking (ADN) charges, and standard data transfer charges are applicable in this design.

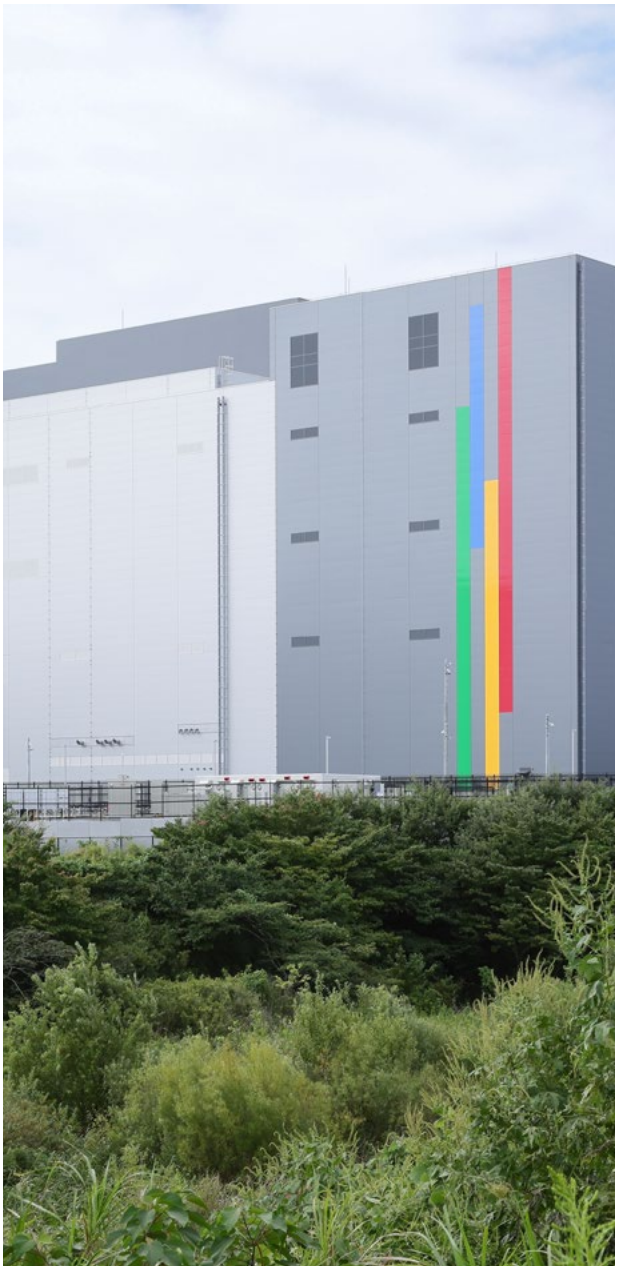

Additionally, customers may need to horizontally scale the number of headends in a single region for high-availability and/or supporting several branch/campus networks as illustrated in Figure 4.2(b). In this case:

- It is essential to also establish overlay tunnels between the individual headends to exchange specific prefixes.
- Each headend can advertise a summary route to the Cloud Router to advertise the aggregate/summarized customer prefixes.
- However, the return traffic could be sent to any of the headends in the given region because Cloud Router cannot discern the source of the aggregate prefix it received from a specific headend.
- The return traffic can be routed to the actual destination using the overlay tunnels and specific prefixes exchanged between the headends.

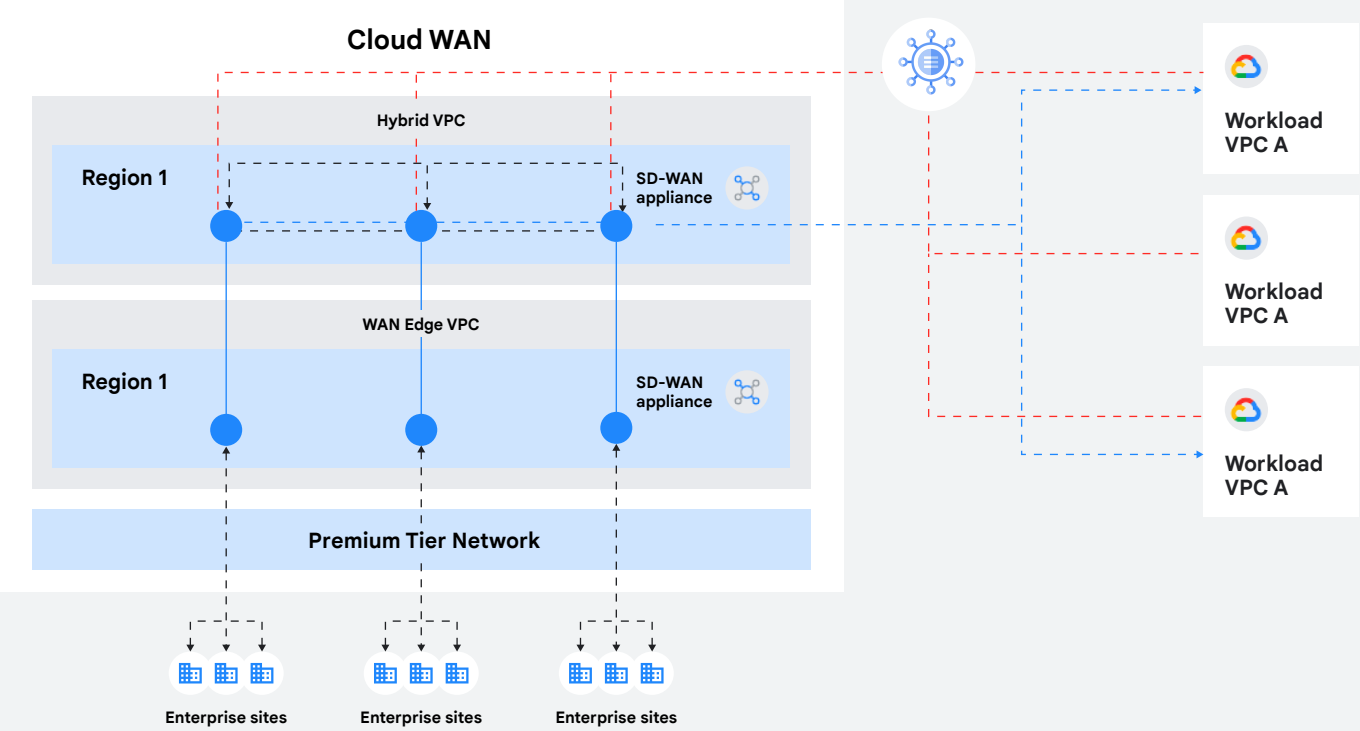

Figure 4.2(b): Horizontally scaled deployment of headends in a single region

Key considerations:

Although Figure 4.2(b) illustrates a single region deployment, the same design can be extrapolated across horizontally scaled deployments in multiple regions with the corresponding overlay network built across all the headends.

c. Branch/campus networks to data centers and/or other cloud providers

This design extends branch/campus connectivity to on-premises data centers and/or other cloud providers. It leverages Google Cloud's Premium Tier network, a multi-NIC NVA, and configures either Cloud Interconnect or Cross-Cloud Interconnect based on the remote destination.

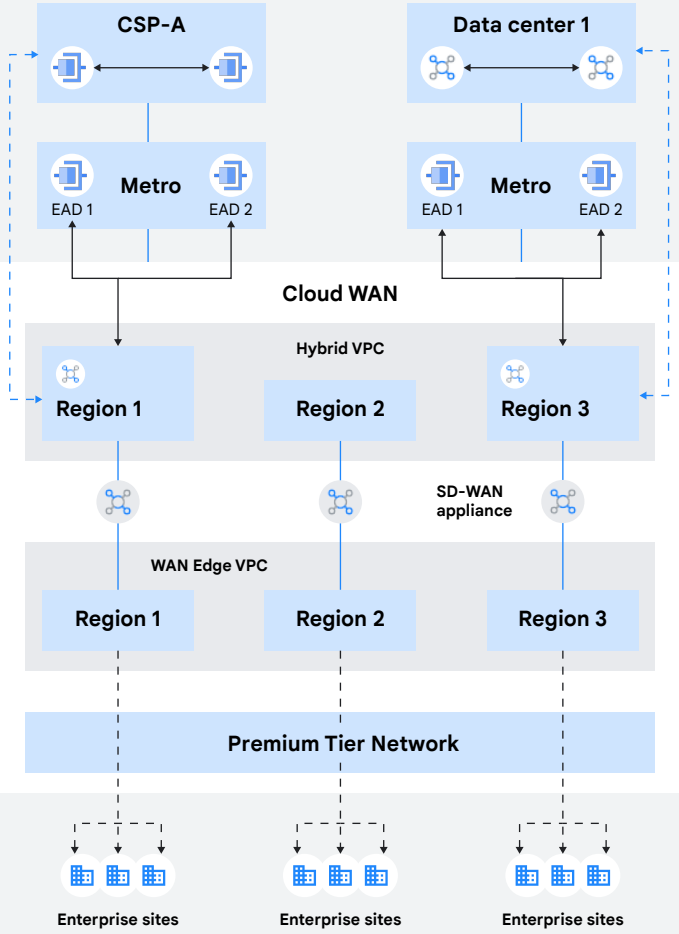

Figure 4.3: Branch/campus networks to a data center and/or other cloud provider

NVA for segregation and routing:
The multi-NIC NVA continues to enable trusted and untrusted network segregation and granular routing. It is deployed across two VPCs:

- WAN Edge VPC: For terminating untrusted (internet) connectivity.
- Hybrid VPC: For configuring private overlay connectivity to data centers or other cloud-provider networks.

Extended connectivity with overlay:
Customers extend connectivity from branch/campus networks via NVAs deployed in a VPC. The BGP peering over the VLAN attachment associated with Cross-Cloud Interconnect or Cloud Interconnect is specifically configured to exchange only the private IP(s) associated with the NVA's secondary virtual network interface within the Hybrid VPC. An additional overlay network is then required to be established directly between the NVA and the corresponding data center or other cloud-provider network to route traffic. This overlay network can be established using IPsec/GRE protocol or proprietary vendor tunneling protocols. This design is particularly recommended when the number of prefixes exchanged exceeds the maximum dynamic route limits of direct NCC advertisement.

Alternative for in-limit prefixes (NCC-based routing): If the number of prefixes remains within the maximum dynamic route limits, NCC can be used to route all branch and site traffic using the underlying VPC network and corresponding NCC spoke attachments. This alternative design, illustrated in Figure 4.4, often eliminates the need for configuring a multi-NIC NVA for routing purposes, as the NVA can be added as a router appliance spoke of the NCC hub to route traffic directly via the underlying VPC network:

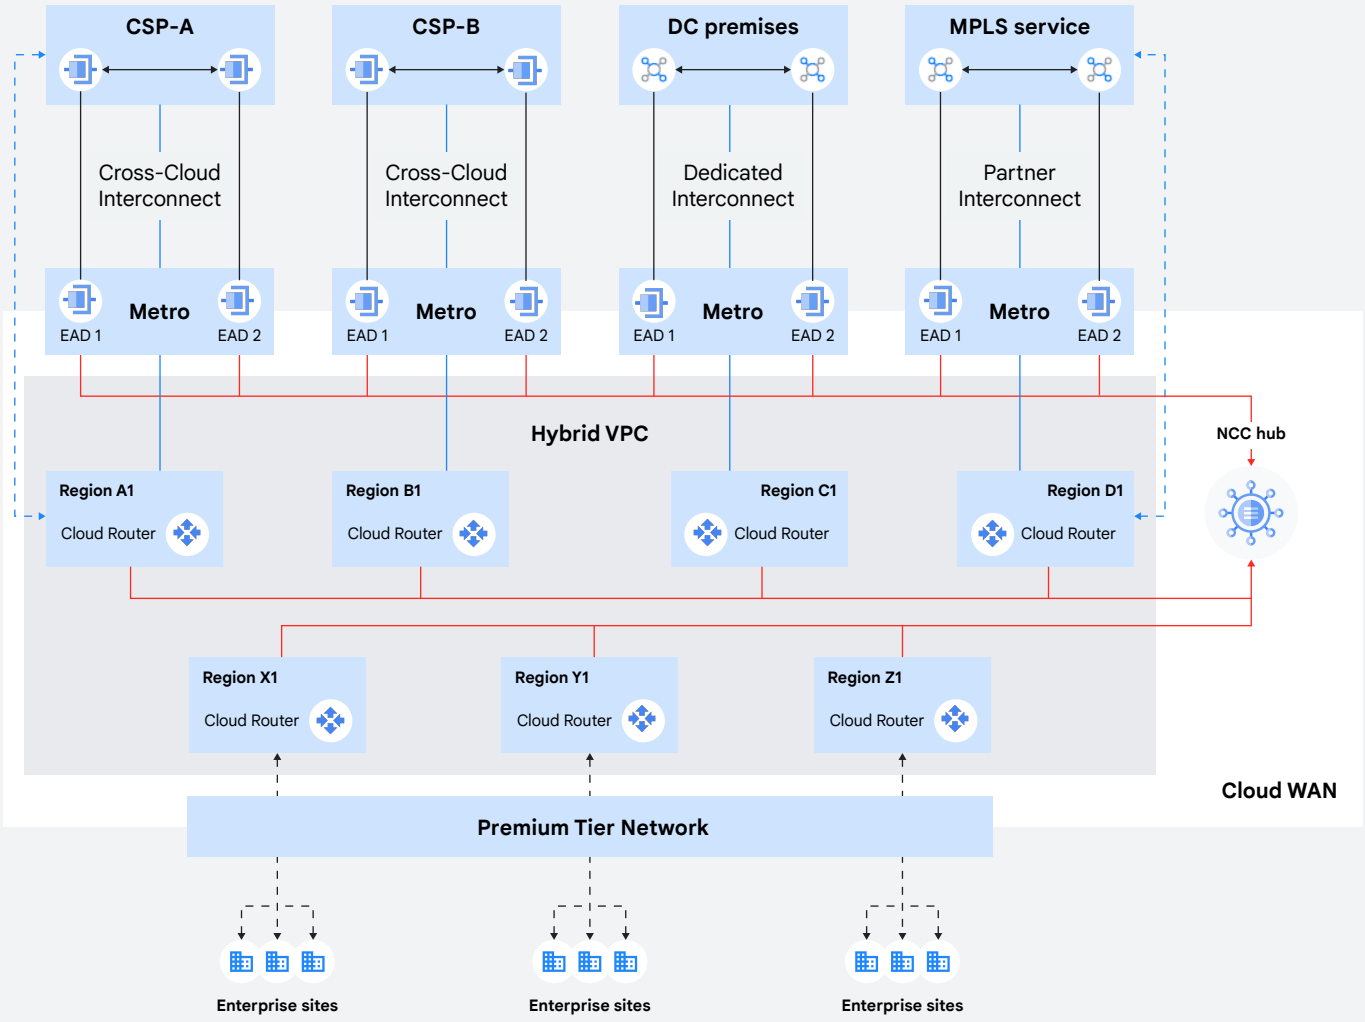

Figure 4.4: Branch/campus networks to datacenter and/or other cloud provider

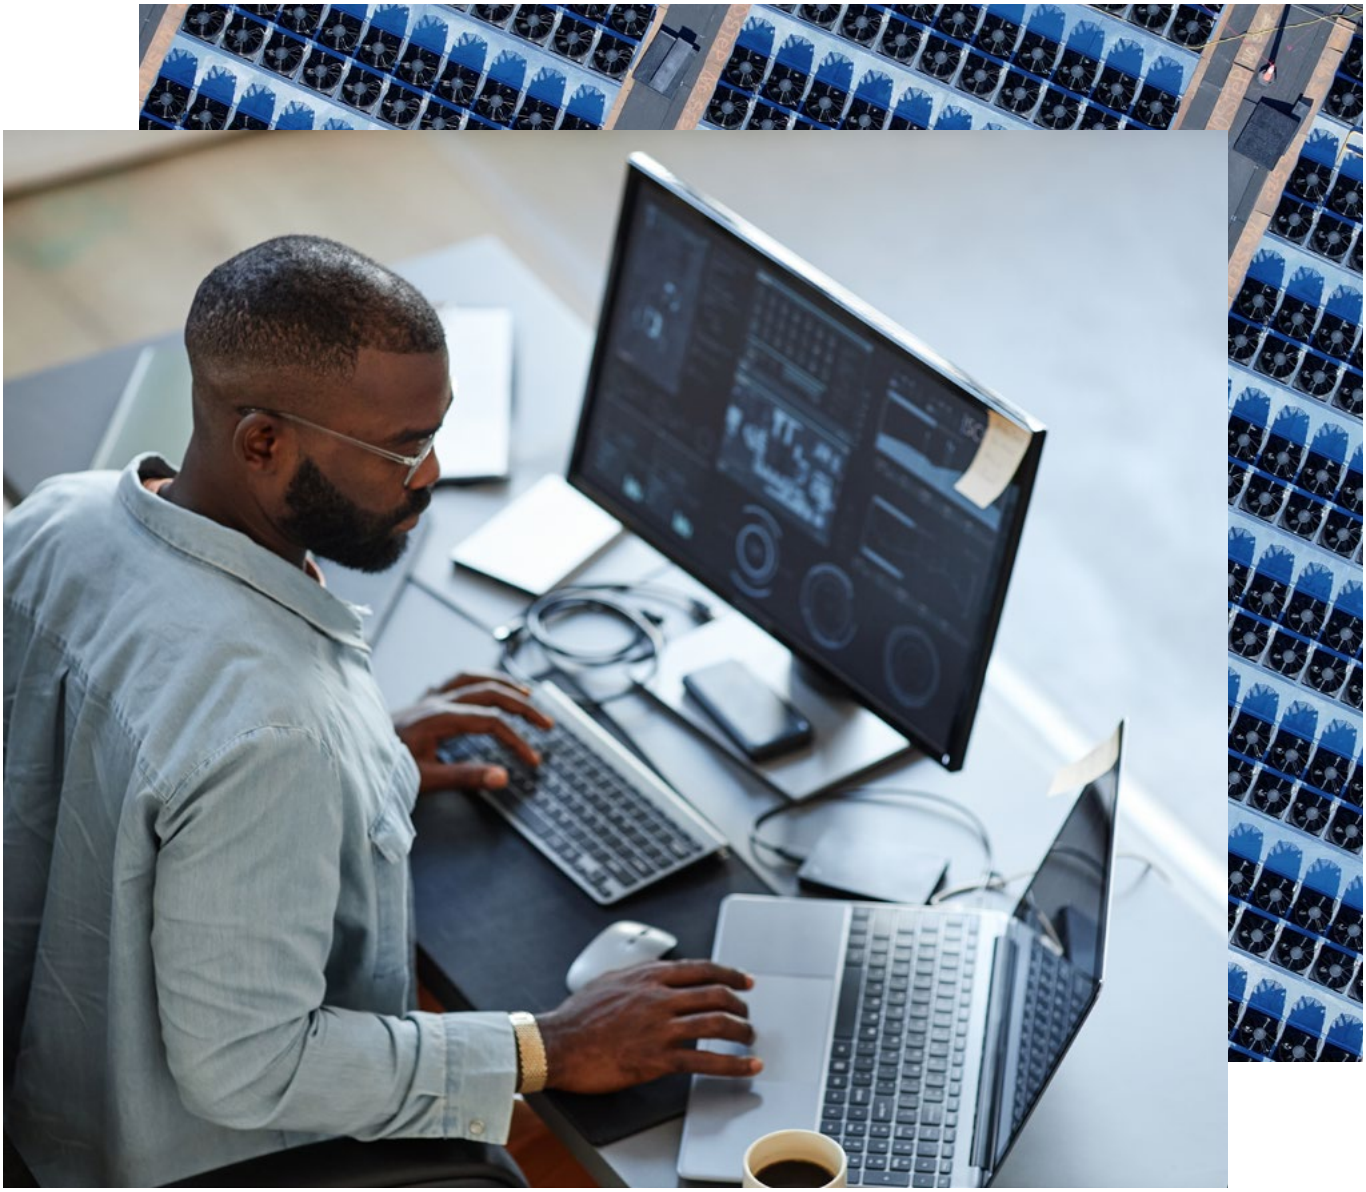

05

Securing workloads on Cloud WAN

Securing workloads connected via Cloud WAN is paramount, especially when customers migrate their branch/campus networks to Google Cloud.

a. Securing private applications

For securing private application traffic originating from branch/ campus networks and destined for workload VPC(s) within Google Cloud, customers can configure Cloud NGFW policies directly on the workload VPCs.

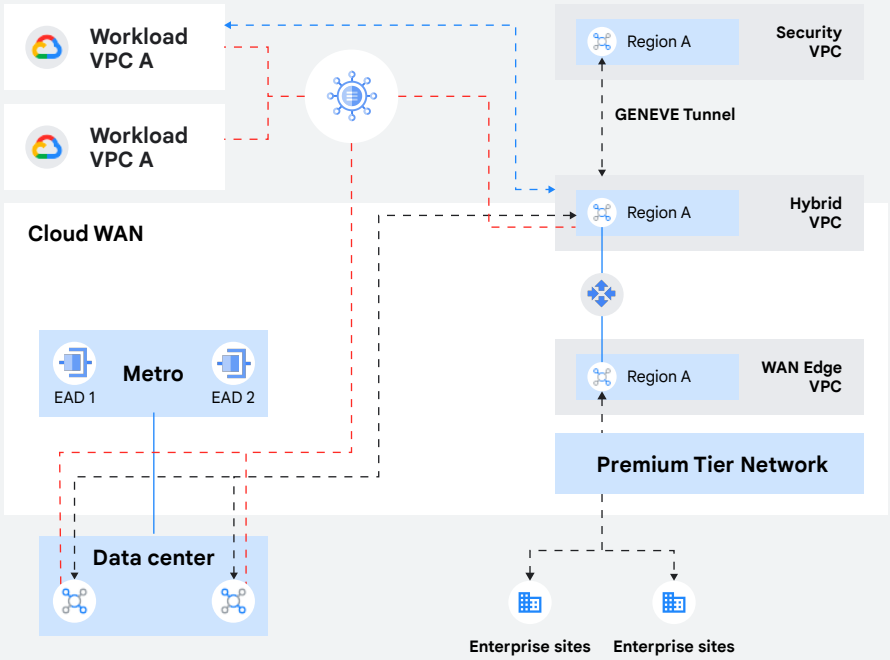

Figure 5.1: Securing private application using In-band NSI

For advanced security capabilities, including intrusion prevention systems (IPS) and intrusion detection systems (IDS), the following next-generation firewall (NGFW) options are available:

Cloud NGFW: Leveraging native Google Cloud Firewall endpoints for policy enforcement.

Third-party NGFWs: Utilizing in-band Network Security Integration (NSI) for seamless traffic redirection to external firewall appliances.

Policy enforcement and traffic interception: The security interception policy is strategically enforced on the trusted network interface card (NIC) of the SD-WAN headend appliance (the interface residing within the Hybrid VPC). This ensures that all traffic terminating on the SD-WAN headend, and intended for the workload VPC(s), is subject to a predefined traffic-matching firewall policy.

Transparent traffic tunneling: Any traffic that matches the configured firewall policy is then transparently tunneled to either the Cloud NGFW endpoints or designated third-party NGFWs via in-band Network Security Integration using packet intercept. This redirection occurs without requiring any complex network reconfigurations.

GENEVE Tunneling for inspection: The design illustrated in Figure 5.1 explicitly leverages packet intercept using In-band NSI to route traffic for inspection. This is achieved by establishing a GENEVE tunnel between the SD-WAN headend and the third-party firewall appliance, ensuring encrypted and encapsulated traffic flow for security inspection.

b. Securing public applications using SSE integration with NCC Gateway

Customers commonly deploy security service edge (SSE) solutions to secure access from branch/campus networks to public applications and SaaS services.

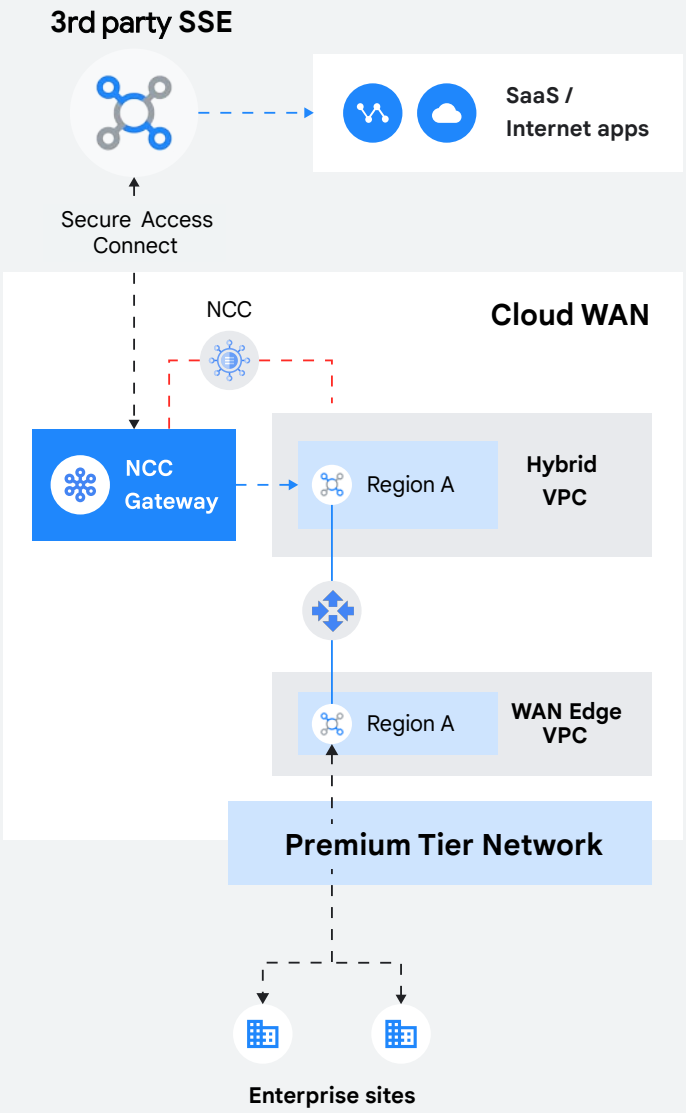

Figure 5.2: Securing access from branch/campus networks to public applications and SaaS services using NCC Gateway

Integrated design with NVA and NCC: This design, illustrated in Figure 5.2, integrates Google Cloud's Premium Tier network, a multi-NIC Network Virtual Appliance (NVA), and a central NCC hub.

NCC Gateway for traffic steering: An NCC Gateway is deployed and added as a dedicated spoke of the NCC hub. This NCC Gateway is engineered to steer aggregated traffic from the SD-WAN appliances to the third-party SSE provider. This architecture facilitates the establishment of private connectivity from the customer's on-premises environment to the SSE provider via Secure Access Connect (SAC).

NVA role in traffic flow: The multi-NIC NVA continues to enable trusted and untrusted network segregation and granular routing. It is deployed across two VPCs:

- **WAN Edge VPC:** For terminating untrusted (internet-bound) connectivity.
- **Hybrid VPC:** For steering traffic towards the NCC Gateway for security inspection by the SSE provider.

Data center traffic security with NCC Gateway: For securing high-performance connectivity originating from a data center and destined for public applications, the Cloud Interconnect is configured to terminate a VLAN attachment directly on the NCC Gateway. This ensures that traffic flows directly to the NCC Gateway, allowing it to be steered to the SSE provider for comprehensive inspection before being routed to the Internet.

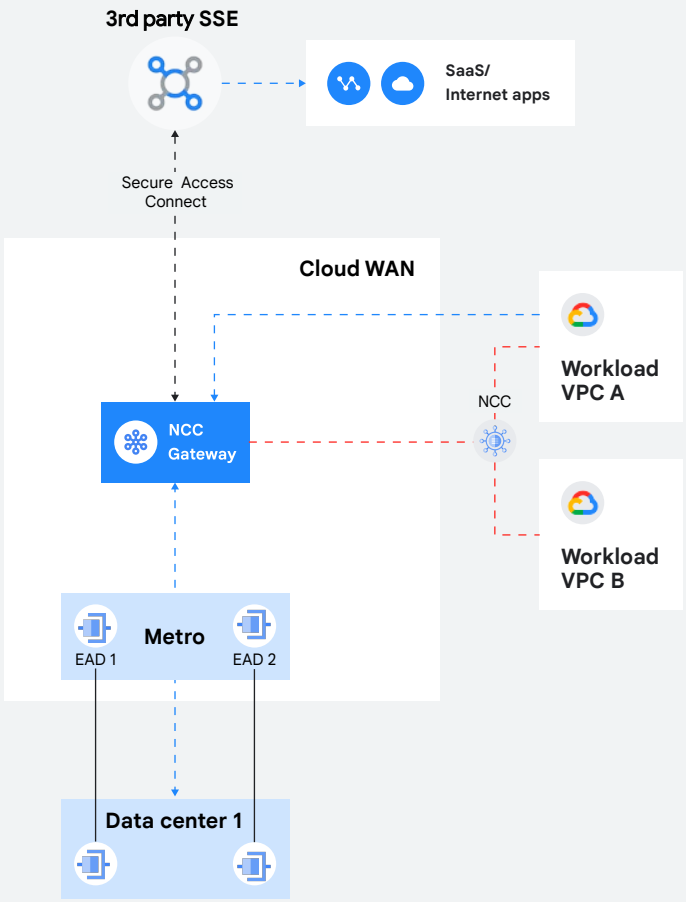

Figure 5.3: Securing access from data centers and workload VPCs to public applications and SaaS services using NCC Gateway

Workload VPC security: Traffic from other Google Cloud workload VPCs accessing public applications can also be secured by adding these VPCs as spokes of the NCC hub. This enables centralized policy enforcement for all public-bound traffic.

Comprehensive security illustration: Figure 5.3 depicts a robust security architecture where traffic sourced from a data center (via Cloud Interconnect) and from workload VPCs (attached as NCC VPC spokes) is secured through the NCC Gateway and SSE integration.

Key considerations:

NCC Gateway is a regional deployment and requires the implementation of a hybrid inspection topology to effectively steer traffic to itself for inspection.

06

Summary and key takeaways

Securing workloads connected via Cloud WAN is paramount, especially when customers migrate their branch/campus networks to Google Cloud.

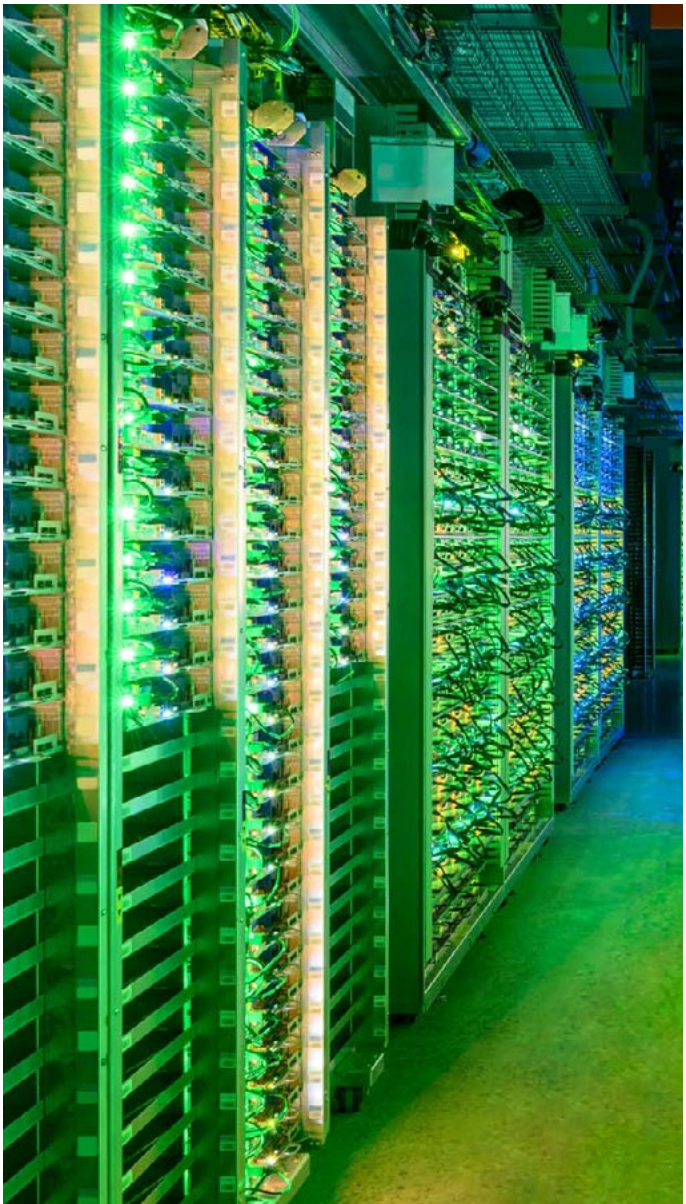

Google's global network

Cloud WAN leverages Google's optimized, global network for enhanced application performance, reliability, and security.

Versatile connectivity

Data center and cloud-to-cloud: Utilize Cloud Interconnect, Cross-Site Interconnect, and Cross-Cloud Interconnect for dedicated private links across on-premises and multicloud environments, with Network Connectivity Center (NCC) as a central hub.

Branch and campus integration: Google Cloud's Premium Tier network, NVAs, and Network Connectivity Center offer scalable options to extend branch/campus networks into Google Cloud for secure resource and SaaS access.

Integrated security

Cloud WAN designs robust security support via Cloud NGFW, third-party NGFWs (in-band Network Service Integration), and security service edge (SSE) integration with NCC Gateway for both private and public application traffic.

Flexibility

These solutions offer architectural flexibility for Layer 2 transparency (Cross-Site Interconnect), centralized routing (Network Connectivity Center), or NVA-based overlays for diverse enterprise needs.

Hybrid and multicloud optimized

These solutions are designed for seamless, high-performance, and secure connectivity in complex hybrid and multicloud environments.


Google Cloud