Google Cloud

# Cross-Cloud Network

Solution Overview

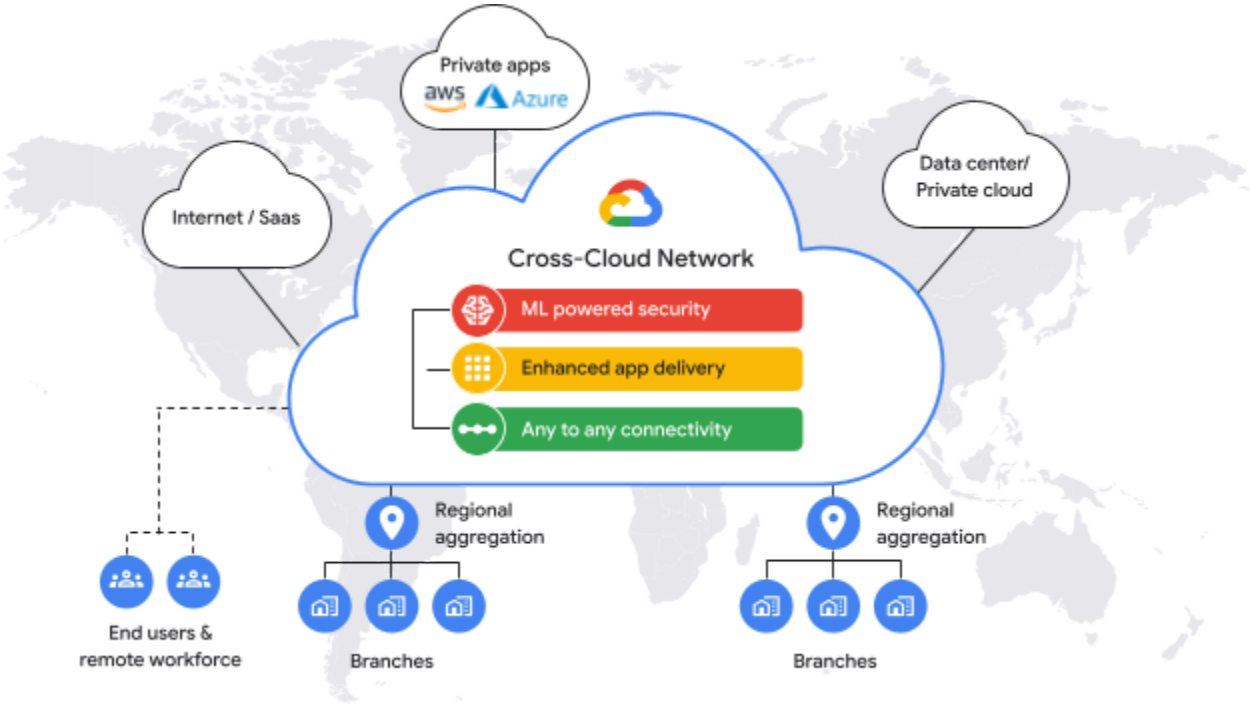

# Optimize multicloud connectivity with Cross-Cloud Network

Connect enterprise networks on Google's global infrastructure for performance, reliability, and security.

### Motivation

Companies optimize applications by sourcing services from multiple cloud providers. This distribution creates an application stack across data centers, imposing stringent privacy, latency, throughput, and security requirements on heterogeneous and disaggregated infrastructure. Simultaneously, both workforce distribution and security stacks have become disaggregated, as office locations connect via secure private networks, while remote users rely on public networks, increasing security vulnerabilities. This trend results in numerous point solutions, requiring businesses to integrate diverse technologies to keep pace with multicloud evolution. Consequently, businesses must engineer custom network backbones to interconnect data centers and users, steering traffic through specialized security stacks. However, the rise of AI-powered applications has made existing multicloud networking solutions inadequate. A new multicloud networking paradigm is essential to achieve the required connectivity and security with the agility, performance, and cost expected of cloud environments. This guide presents a Cross-Cloud Network solution designed to address these critical challenges.

### Solution

Cross-Cloud Network is a global cloud networking solution with any-to-any connectivity, enhanced application experience and ML-powered security across all users and workloads wherever they may be. Cross-Cloud Network provides a fabric to interconnect the different networks in a global multicloud environment and to optimally deploy security services, making it the hub for connectivity, security, and application delivery services in a multicloud environment. By centralizing these services through Cross-Cloud Network, enterprises can simplify the connectivity and security challenges of supporting multicloud applications, while optimizing the connectivity paths and performance that the applications require. Cross-Cloud Network addresses the multicloud challenges by delivering innovative services to enhance application delivery across multiple clouds with cloud-native ML-powered security stacks over Google's planet-scale network to deliver pervasive connectivity with performance, reliability and ubiquitous presence.

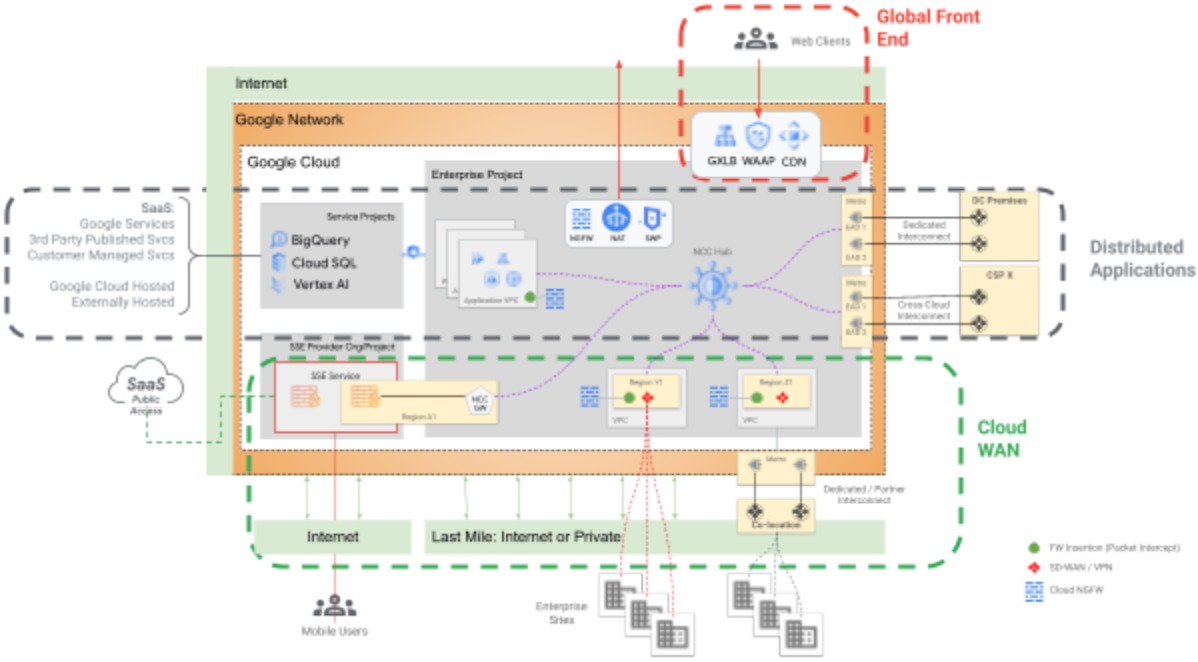


By bringing this functionality, Cross-Cloud Network optimally addresses the following use cases:

- **Distributed applications** to efficiently compose applications leveraging best of breed services hosted across cloud providers and on-prem
- **Global front end** to securely and efficiently serve web origins distributed across a hybrid multicloud infrastructure
- **Cloud WAN** to enable secure connectivity for the hybrid workforce by leveraging the global scale, reliability, and ubiquity of the Google network.


### Planet scale any-to-any connectivity

Having a global presence can be challenging for enterprises as vendors, regulations, operations, and pricing vary significantly across geographies. The distributed nature of business will often call for highly distributed connectivity in order to optimize the application experience and associated productivity. Cross-Cloud Network enables the elastic use of Google's global network, to give the enterprise global presence, cloud-grade performance, and reliability. With 202 Google Cloud points of presence (PoPs) distributed across over 200 countries and territories worldwide, customers can use Cross-Cloud Network to interconnect cloud networks and private locations elastically at any location.

Cross-Cloud Interconnect streamlines connectivity between cloud providers by having Google Cloud handle the provisioning of multicloud high-speed connections on behalf of the user. Cross-Cloud

Interconnect provides a secure, high-speed direct connection between cloud service providers to guarantee privacy, throughput, availability, and predictable latency. By using Cross-Cloud Interconnect, the customer does not need to maintain any infrastructure in co-locations or on-premises. Private locations are easy to connect to the Cross-Cloud Network using any of the hybrid connectivity options available, which include: IPsec-based Cloud VPN tunnels, Partner Interconnect, Dedicated Interconnect, or your choice of SD-WAN solution enabled by the Network Connectivity Center.

**Build distributed applications and enhance application delivery**

As applications modernize and migrate to the cloud, dependencies and preferences on specific cloud services develop. As a result, different applications will be hosted in different cloud providers. Furthermore, some applications may be built with a combination of resources and services that are distributed across multiple clouds. This may be generally seen as an exercise in calling APIs across the different clouds. Before these APIs can be called, proper connectivity channels must be established between the applications and their component resources and services so that every call doesn't travel out to the internet and back into the cloud, which would be both insecure and inefficient.

Cross-Cloud Network enables multicloud communication between application components over optimal private connections by extending the capabilities used in Google Cloud to resources and services in other cloud providers. Cloud Load Balancing provides mechanisms to group backend compute resources, balance the workload across the group members, and achieve elastic capacity for different application components. The group of resources is reachable via the front-end IP of the load balancer, effectively encapsulating the group of resources as a single resource IP that can be used to assemble the application. Global access to the load balancer front-ends allows the flexibility to include resource groups from any region in an application.

With global backends, backend resources may be distributed across regions, enabling multi-region resiliency and optimal load distribution across geographies. With the introduction of hybrid network endpoint groups, the backend resources can be reachable across a hybrid connection such as Cloud Interconnect or Cloud VPN. The functionality includes health checks that stretch into other networks over the hybrid connection. This effectively allows the user to model resources and services hosted in other clouds as IP end-points that are natively reachable in the Google Cloud. When the resource group needs to be consumed as an API, this abstraction can be refined further by using Private Service Connect to provide a private address for the resource group that can be accessed in the private network, across regions and/or over a hybrid connection. By having these abstractions, resources from disparate clouds and under disparate connectivity models can be brought together under a unified connectivity model that enables the effective assembly of an application across a multicloud environment.

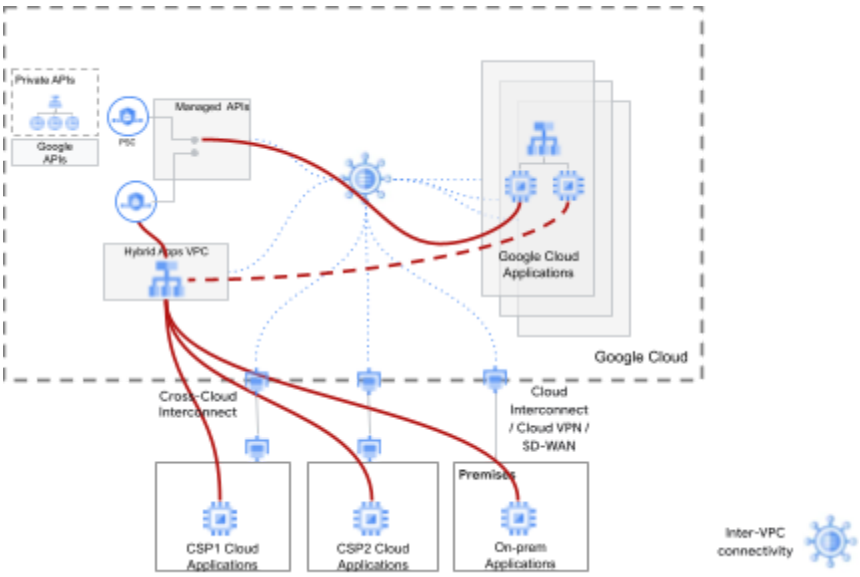


As we assemble applications in a cloud-native manner and normalize the model to that used in Google Cloud, applying security between the application components is now easier. Cloud Next Generation Firewall (NGFW) can be used to apply firewall-level inspection to the application flows that happen within Google Cloud and also with other clouds and on-premises data centers. The in-band insertion of the firewalls is enabled by packet interception using traffic matching policies at the workload interfaces controls is cloud-native, the enforcement of the policies is done by firewall engines provided by our ecosystem partners. An NGFW service may be fully managed by Google Cloud with Cloud NGFW, in which case the policies are authored as an extension to the firewall policies in the VPCs, or it can be a self-managed 3rd party NGFW inserted using in-band Network Security Integration (NSI) orchestrated by the Network Services Integration manager. When integrating 3rd party NGFWs with NSI, users can access the vendor-specific policy console for policy authoring.

**Deliver internet-facing web applications with the Global Front End**

Google Cloud's global load balancer provides the opportunity for users to unify their Global Front End. As part of the Global Front End, Google Cloud Armor provides a web application firewall (WAF) and distributed denial-of-service (DDoS) mitigation service. It helps protect websites and services hosted in any cloud from multiple types of threats, including:

- DDoS protection: Cloud Armor can mitigate both volumetric and targeted DDoS attacks.
- WAF protection: Cloud Armor can block a wide range of web application attacks, such as XSS, SQL injection, and denial of service attacks.
- Bot management: Cloud Armor can help you identify and block bots, like scrapers and spammers.

The global external Application Load Balancer offers an ecosystem of Service Extensions powered by an Envoy proxy infrastructure. Service Extensions are instrumental in delivering API security services and other protection that is critical to the Global Front End.

For public application access, the backend behind the global external Application Load Balancer can include a mix of resources hosted in Google Cloud, on-premises or other cloud providers. This effectively allows customers to centralize their Global Front End for public facing applications with

Cross-Cloud Network, rather than maintaining different front ends with disparate models in different clouds. The Global Front End natively balances connections based on the geographic distribution of the demand to minimize latency and improve application experience. Rather than relying on elaborate DNS localization schemes, the global external Application Load Balancer offers a global anycast IP front-end coupled with connection localization. Traffic is sent to the global anycast IP from anywhere in the world and is promptly connected to the nearest instance of the global load balancer, which steers the connection to the closest relevant backend resources. Thus, connections are seamlessly optimized on a global basis for speed and latency. Also, applications using Google Cloud's Global Front End will have access to Cloud CDN to cache high demand static content and again improve application performance and user experience.

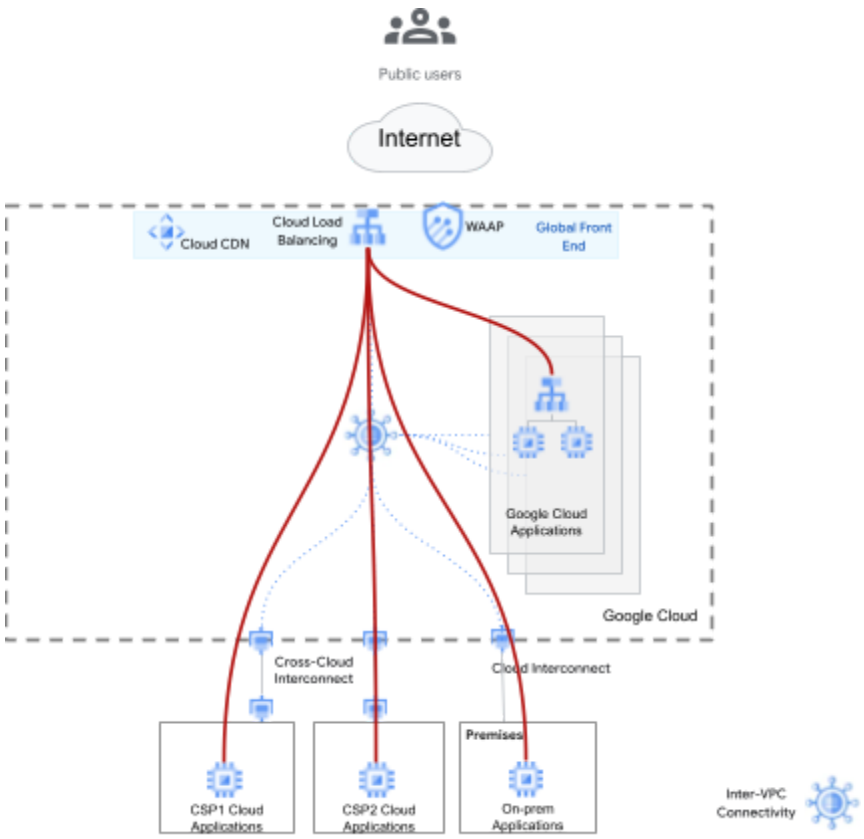


### Secure employee access to applications with Cloud WAN

As applications move to the cloud and the workforce turns to a hybrid model between the office and remote locations, the enterprise ends up managing a multitude of disjointed security stacks. Not only is this risky and complex, but it also requires serious compromises in the performance of the connectivity as traffic takes suboptimal detours in order to be inspected by different security stacks. With Cross-Cloud Network, connectivity and the delivery model for the application are consolidated, paving the way for the consolidation of the security stack and the start of the journey to make these security services truly cloud-native.

Employees are a key vector for security attacks. As employees adopt a hybrid workstyle, the risk is even greater. This risk has been managed by procuring a security stack that is application aware, assisted by artificial intelligence to recognize patterns and anomalies and rich in user and end-point authentication controls. Employees connecting over a public network to any resource must traverse such a security stack, the current industry term for such a stack is the security services edge (SSE). Beyond the threat

and malware protections that an NGFW provides, SSE stacks include cloud access security broker (CASB), data loss prevention (DLP), and other identity verification services critical to securing connections over public networks like the internet.

There are many providers of SSE stacks, and most offer the stack as a managed service. When offered as a managed service, the SSE stack is hosted at a series of locations around the world and reachable over the internet. To maintain secure communications, traffic between the SSE stack and the applications must be encrypted and tunneled, reducing the effective throughput of the security stack. Whether employees are connecting over the internet or working from a company location, traffic must be steered through a security stack. Many of the SSE services pertinent to the public nature of the internet may not be required when users connect over private links from company locations. This has resulted in enterprises managing separate stacks to secure employees working from the office vs. employees working outside the office. This presents a challenging combination of network complexity, performance, and disparate security stacks.

Cross-Cloud Network provides the necessary functionality to steer company user traffic to a security stack that is hosted in Google Cloud. The security stack can be a fully-managed SSE stack provided by Google's ecosystem partners, or it can be a more focused NGFW stack delivered as a Google-managed or a 3rd party cloud-native service, or a combination of these options. Traffic can be steered according to policy to the appropriate security stack, making the routing to the SSE or NGFW seamless. Since the security stack is deployed natively in Google Cloud, there are no tunnels or encryption required, effectively bringing the performance of the stack to its full potential. The security stack also enjoys the benefits of elastic capacity as it is organized as a producer service with an elastic backend behind the cloud-native load balancers. As traffic is centralized, the security stack can be consolidated. With Cross-Cloud Network, employee security controls are more robust, simpler, and don't impact application performance.