

Google Cloud Cross-Cloud Network Solution Brief

Modern networking for hybrid
and multi cloud enterprise


Preface



Enterprises are leveraging cloud to modernize their infrastructure for business agility.


Today, a majority of enterprises operate in hybrid and multi cloud environments, and streamlining the network has become more important than ever as applications and data become distributed, cybersecurity threats such as bot attacks are skyrocketing, hybrid workforce become table stakes, and AI/ML is leading a new inflection point in business efficiency.

Cross-Cloud Network addresses many of the foundational networking and security needs to enable this new enterprise paradigm. We simplified cloud networking with global VPCs, Private Service Connect, Network Connectivity Center, and Cross-Cloud Interconnect. We are excited to introduce Cross-Cloud Network with use case based solutions and invite customers and partners to join us on this journey.



Introduction


In the quest to leverage a curated mix of cloud application services, the diverse components of cloud applications are increasingly being distributed across a multitude of cloud and on-prem data centers.



Companies seek to utilize the services they consider to be optimal for their applications and must therefore source these services from multiple cloud providers; for instance, an application may use AI services from Google Cloud, while maintaining business critical data on-premises and leveraging other application services from another cloud provider.

This results in an application stack that is distributed across multiple data centers; communications between these components pose stringent privacy, latency, throughput and security requirements on an infrastructure that is inherently heterogeneous and disaggregated.

¹[What Are Enterprise Multicloud Adoption Trends](#), Andrew Smith, IDC #US48902122 (March 2022)



In recent years, the workforce has grown accustomed to a hybrid schedule in which work is done from a combination of private office locations and remote locations dispersed around the world.

Office locations are connected over a network that offers privacy and a certain level of security assurances, but as users move outside of the office, they connect over public networks which create a higher exposure to security threats. The business requires different levels of security for these different connection methods. Hence, not only has the workforce become more distributed, but the security stacks necessary to secure the workforce as they connect over public or private networks have become disaggregated.

The dispersion of applications and consumers across disparate and independent networks creates a significant challenge for businesses that find themselves having to engineer custom network backbones to interconnect the different data centers (on-prem and multicloud) and users; while steering different types of traffic through a variety of security stacks, each of which specializes in addressing the risk profile associated to a specific type of connection.



As this trend evolves, a multitude of point solutions to specific issues emerge in the market and businesses find themselves tasked with integrating these diverse solutions in an attempt to keep up with the pace of the multicloud evolution.

The result is a very complex and costly network infrastructure, which is inevitably sub-optimal and requires the business to coordinate a multitude of providers, technologies, business relationships and also train the operations personnel in the nuances of the different cloud environments.



This situation has been somewhat tenable until now. However, the prevalence of AI-powered applications and Gen AI has created an inflection point in which multicloud networking has become critical and can no longer be realized with the existing toolset and models. A new multicloud networking paradigm is needed to achieve the required connectivity and security with the right agility, performance, and cost to match the cloud standard.



Introducing Cross-Cloud Network


The Cross-Cloud Network is a global open networking platform with any-to-any connectivity, enhanced application experience, and ML-powered security across all users and workloads wherever they may be.

The Cross-Cloud Network provides a fabric to interconnect the different networks in a global hybrid and multicloud environment and optimally deploy security services. It is therefore the hub for connectivity, security, and application delivery services in a hybrid and multicloud environment.


By centralizing these services through the 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.

Google Cloud


Cross-Cloud Network is built upon Google's state-of-the-art fiber optic network that covers the globe. It provides a full stack of intelligent network services that optimizes performance, efficiency, and operational costs to help enterprises run on the most modern infrastructure on the planet. Cross-Cloud Network provides the planet scale any-to-any connectivity, enhanced application delivery, and cloud native ML-powered security.



**Planet scale
any-to-any
connectivity**



**Enhanced
application
delivery**



**Cloud native
ML-powered
security**

Cross-Cloud Network leverages best-in-class products from Google Cloud and partners, enabling enterprises to use cloud-native and partner appliances and services to build a seamless and secure network fabric. Customers have the choice of bringing their network virtual appliances or using Google Cloud services. In the next sections, we will provide an overview and the key use cases of Cross-Cloud Network.




Planet Scale any-to-any Connectivity

Having a global presence can be challenging for enterprises as vendors, regulation, operations, and pricing vary significantly across geographies, yet the distributed nature of business will often call for highly distributed connectivity in order to optimize the application experience and associated productivity.

The Cross-Cloud Network enables the elastic use of the Google Cloud network, to give the enterprise global presence, cloud grade performance, and reliability.



Google Cloud has a presence in over 200 countries and territories. With 187 Google Cloud Points of Presence (POPs) distributed worldwide, customers can use the Cross-Cloud Network to interconnect cloud networks and private locations elastically at any location.



- Network
- Current region with 3 zones
- Future region with 3 zones
- Edge point of presence

Google Cloud introduced Cross-Cloud Interconnect to streamline the connectivity between cloud providers by having Google Cloud handle the provisioning of multicloud high speed connections on behalf of the customer.

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, and it is backed by Google Cloud with a 99.99% SLA.

Cross-Cloud Interconnect supports direct connectivity to major cloud providers such as AWS, Azure, Oracle, and Alibaba Cloud. Private locations are easy to connect to the Cross-Cloud Network using any of the hybrid connectivity options available, which include: IPsec based HA-VPN tunnels, Partner Cloud Interconnect, Dedicated Cloud Interconnect, or your choice of SD-WAN solution enabled by Network Connectivity Center.

Enhanced 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, and 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, yet 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.

The 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.

Google 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. Furthermore, with global backends, 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 an HA-VPN with health checks that stretch into other networks over these hybrid connections.

This effectively allows the user to model resources and services hosted in other clouds as IP endpoints 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 (PSC) to provide a private address for the resource group that can be accessed in the private network, across regions, across organizations and/or over a hybrid connection that may span multiple providers.



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.




When APIs are published using PSC, developers have access to traffic management functionality that enables canary deployments as well as service rate limiting to support the different phases of their CI/CD pipelines.

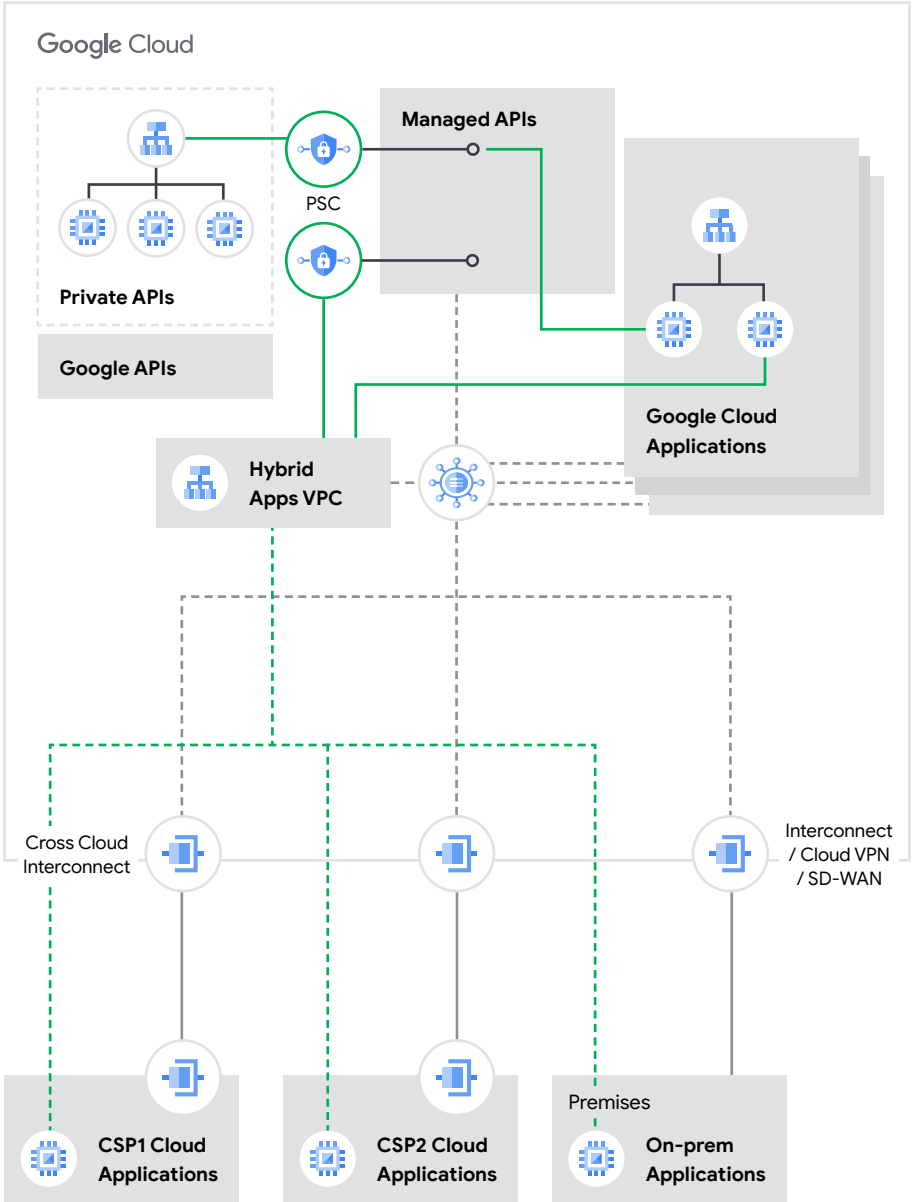


Private Service Connect coupled with load balancing functionality enables a service centric approach to cloud networking. Service centricity in the cloud allows the DevOps, NetOps and SecOps teams to focus on their responsibilities while supporting each other.

DevOps teams can leverage PSC to publish services to an agreed upon landing point in the network topology; NetOps teams can focus on designing the network for reliable, optimal connectivity to the prescribed landing point; and the SecOps teams benefit from a consolidated policy enforcement surface and consistent policy and posture with identity based mTLS zero trust security across workloads, services and data.

ML-powered malware and threat prevention, coupled with in-transit data loss prevention (DLP) ensures that all possible application connectivity paths are secured.





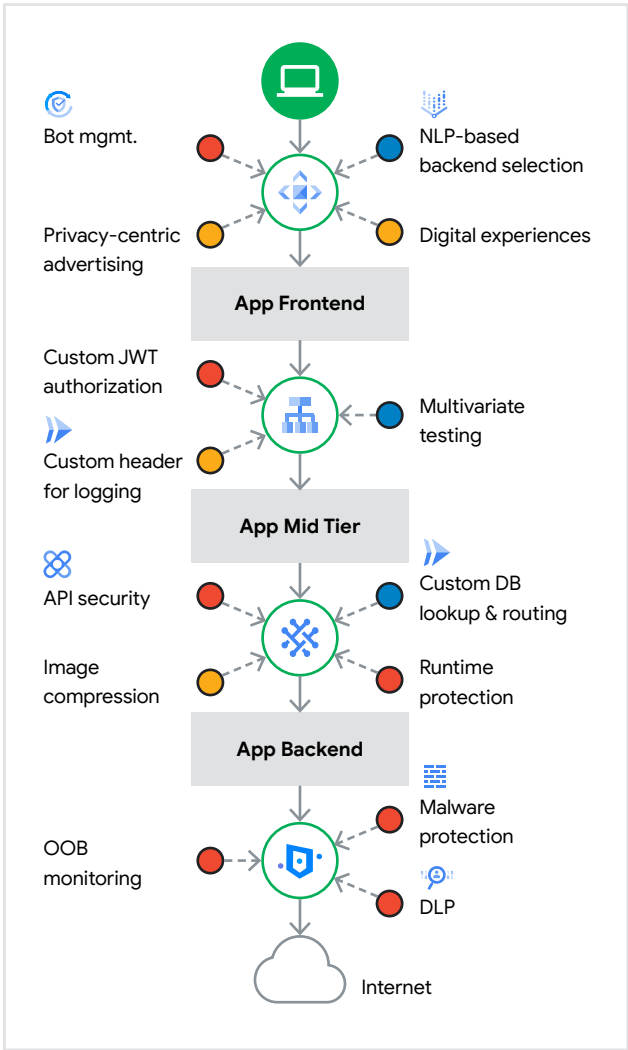
 Inter-VPC connectivity (NCC or Classic Peering)

--- Network connections

- - - Hybrid NEGs

— Application Services connectivity

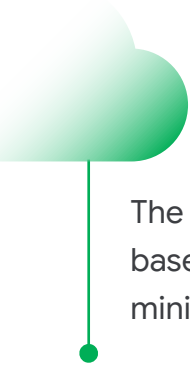
The PSC and load balancing components of the service centric network also enable the insertion of service extensions that enable the seamless integration of routing, security and traffic functionality across the application stack. Service extensions make the Cross-Cloud Network extensible and open by enabling the seamless integration of an open ecosystem of partner solutions.



Google Cloud

Similarly, for public application access, the backend behind the Google Cloud global load balancer (in the Google Cloud global front end) 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 the 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 Google Cloud global 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.

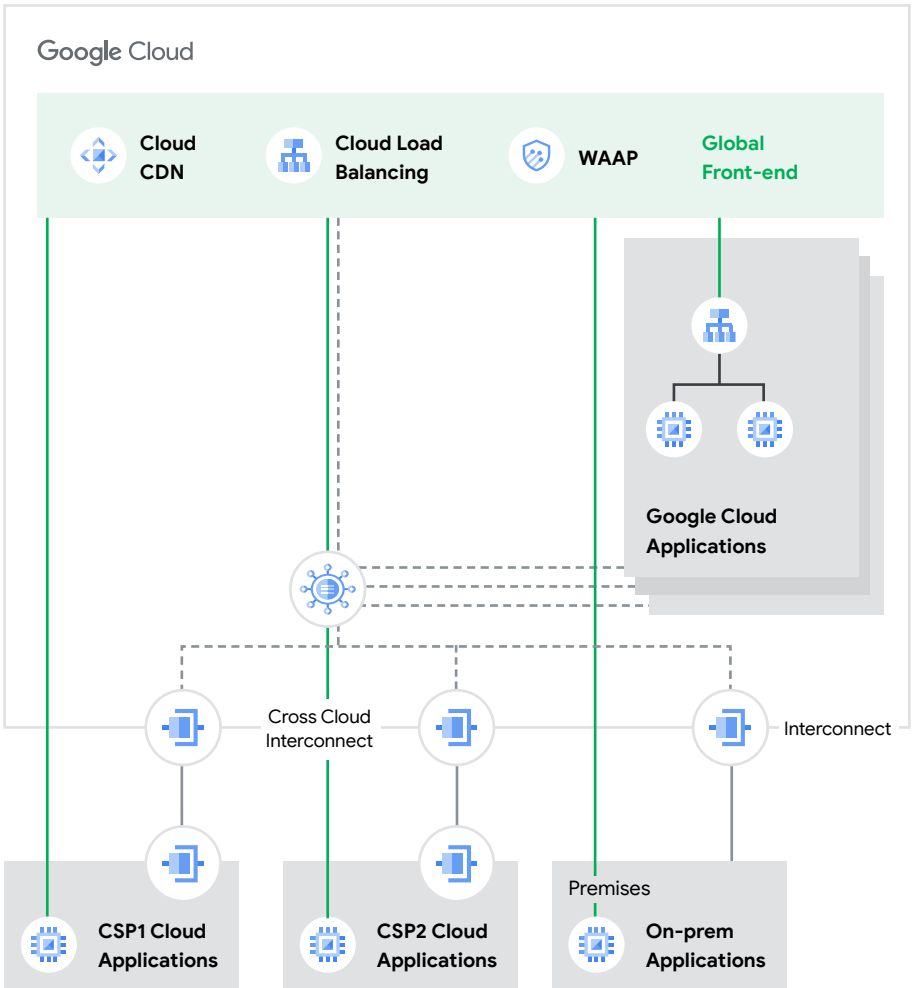
Thus, connections are seamlessly optimized on a global basis for speed and latency. Also, applications using the Google Cloud global front end will have access to Google Cloud CDN to cache high demand static content and again improve application performance and user experience.



Public users



Internet



Inter-VPC connectivity (NCC or Classic Peering)

--- Network connections

— Load Balancer backend connections



Cloud Native ML-powered Security Services

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 disjoint security stacks.

Not only is this risky and complex, but it also requires serious compromises in the performance of the connectivity as traffic takes sub-optimal 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.



The Cross-Cloud Network takes a holistic approach to cloud native security and introduces enhanced security in the following areas:

Workload security: Identity and Access Management (IAM) based zero trust access to workloads and APIs. Malware and threat prevention with cloud-native integration of next generation firewall (NGFW) stacks, and distributed denial-of-service (DDoS) attack prevention for web origins with Cloud Armor.

Data security: Data leak prevention for data at rest and in-transit are integrated into Cross-Cloud Network.

User security: Cloud native security service edge (SSE) managed service integration and Chrome Secure Enterprise browser provide secure user connectivity from any location, on any device to any application.

There are three types of connection to be secured:



Public access to public applications



Employee access to applications (public and private)



Application to application communications

Public access to public applications

As previously mentioned, Google Cloud's global load balancer provides the opportunity for users to unify the global front end, but connections through this global front end must also be secured. Google Cloud Armor provides a web application firewall (WAF) and DDoS mitigation service to help 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 Application Load Balancer offers an ecosystem of Service Extensions powered by an envoy proxy infrastructure. These Service Extensions are instrumental in delivering API security services and other protection critical to the global front end.



Employee access to applications

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 service edge (SSE). Beyond the threat and malware protections that an NGFW provides, SSE stacks include cloud access service broker (CASB), 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, in order to maintain communications secure, traffic between the SSE stack and the applications must be encrypted and tunneled, reducing the effective throughput of the security stack significantly.

Whether employees are connecting over the internet or they are 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. The challenge is therefore a combination of network complexity, performance, and disparate security stacks.

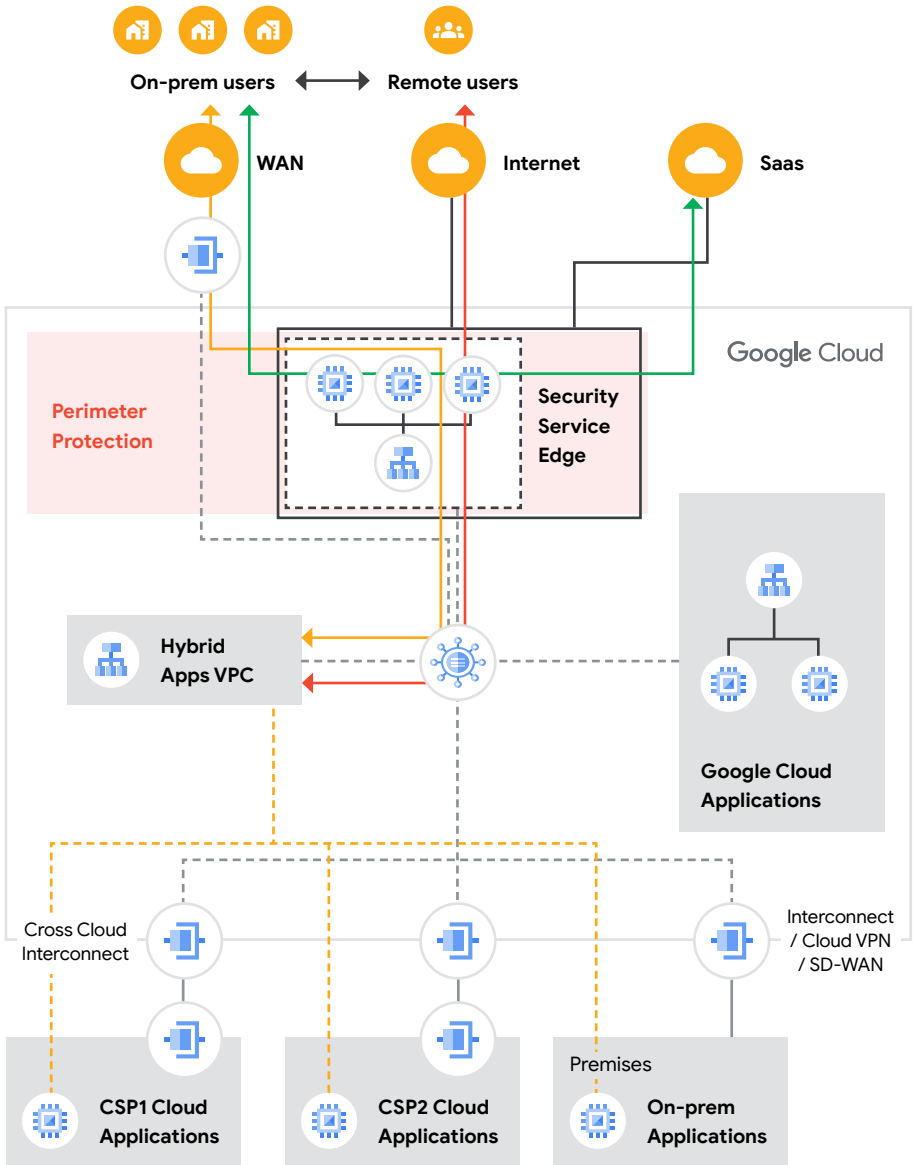
The Cross-Cloud Network provides the necessary functionality to steer traffic for company users to a security stack that is hosted in Google Cloud. The security stack can be a fully managed SSE stack provided by Google's ecosystem partners², or it can be a more focused NGFW stack delivered as a first or third-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 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.

Secure access services edge (SASE) solutions use VPNs for the connection of mobile users. Enterprises have to manage a large number of agents to implement these VPNs. Managing these agents is inconvenient and does not address any use of unmanaged devices by the workforce.

The Cross-Cloud Network (CCN) integrates with Chrome Enterprise Premium to offer a Secure Enterprise Browser (SEB) to secure workforce communications from managed and unmanaged devices.

² Fortinet, Palo Alto Networks, Symantec



 Inter-VPC connectivity (NCC or Classic Peering)

— Remote user to app communication

— On-prem user to app communication

— User to SaaS communication

--- Network connections

- - - Hybrid NEGs

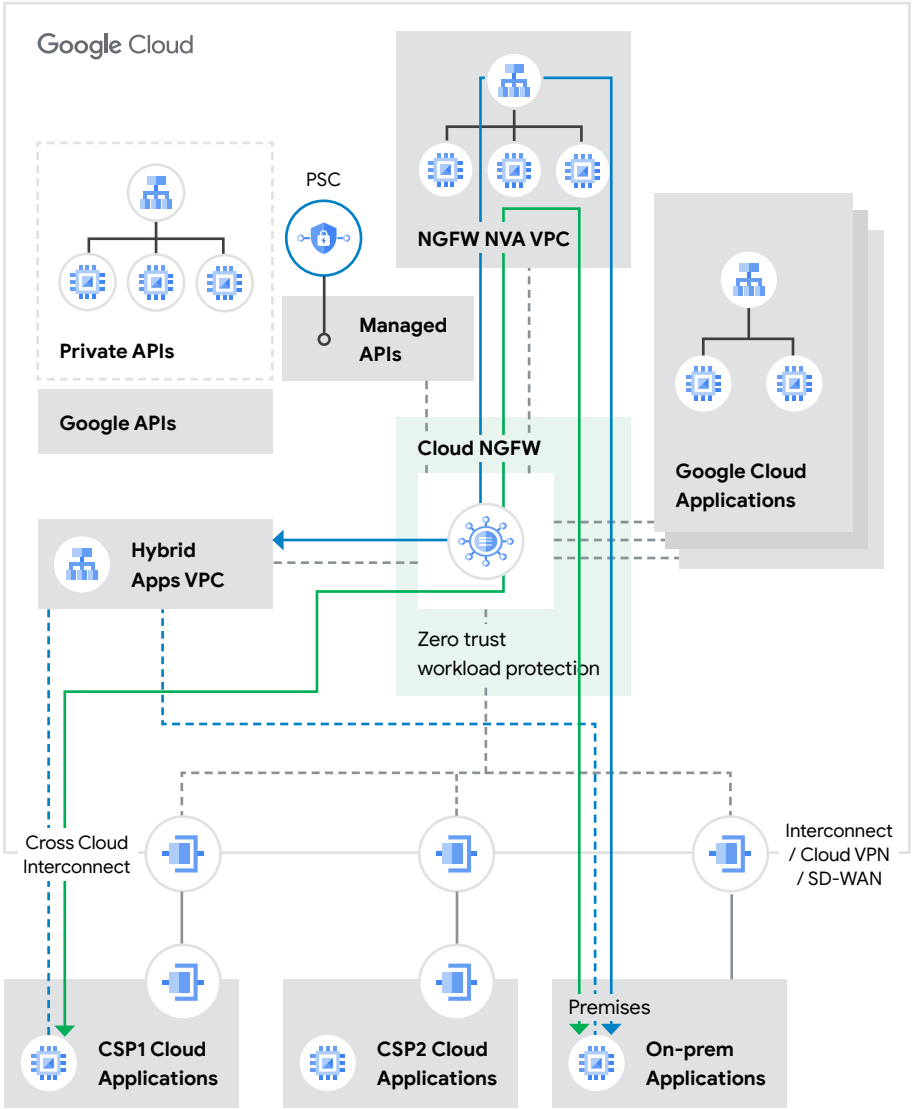
Application to application communications







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 NGFW can be used to apply next generation firewall level inspection to the application flows that happen within Google Cloud and also with other clouds and on-premises data centers.

The insertion of the firewall controls is cloud native, the enforcement of the policies is done by firewall engines provided by our ecosystem partners.

The NGFW service may be fully managed (Google Cloud NGFW), in which case the policies are authored as an extension to the Cloud Firewall policies in the VPCs; or it can be a self-managed NGFW, in which case the deployment is orchestrated by the Network Services Integration Manager and users can access the vendor specific policy console for authoring of the policy.

For full coverage of all possible connectivity flows, a third party NGFW network virtual appliance (NVA) can be inserted in the traffic path leveraging advanced routing functionality in Google Cloud. Policy based routing (PBR) enhances the functionality of the VPC routing stack to allow for policy based insertion of the security stack without the modification of the VPC connectivity design.



-  Inter-VPC connectivity (NCC or Classic Peering)
-  Transit traffic w/Perimeter Security
-  Transit traffic w/Perimeter Security Load Balanced
-  Workload protection with Cloud NGFW
-  Network connections
-  Hybrid NEGs

Workload security is hardened by enforcing Identity and Access Management (IAM) based policies. IAM based security policies in CCN leverage Certificate Authority Authorization (CAA) and mutual-TLS to evolve cloud workload security to a hardened zero trust model.

Cloud Data Loss Prevention (DLP) is instrumental in preventing exfiltration of data at rest. CCN expands DLP protections to data in transit. By leveraging Service Extensions in the Cloud Load Balancers and Google Cloud [Secure Web Proxy](#) (SWP) leading DLP in transit solutions from technology partners³ can be applied to workload and internet egress traffic.

Key Use Cases of Cross-Cloud Network



Cross-Cloud Network enables many use cases today and provides a foundation for any company who may have hybrid or multicloud resources in the future.

We will focus on three initial uses cases which are:



Distributed applications



Internet-facing application and content delivery



Hybrid workforce


³ Broadcom

Distributed Applications

A cross-cloud distributed application is an application stack that combines services from multiple clouds and private data centers. In contrast with the traditional use of multicloud networking, the Cross-Cloud Network is built to enable the consumption of preferred services across multiple clouds and their combination into an application stack that is effectively distributed across multiple clouds and private data centers.

Traditional multicloud networks are built to enable consumption of different application stacks hosted in different clouds with all layers of an application stack hosted in a single cloud provider or private data center. Multicloud connectivity can then be used for one cloud to backup another, or to transfer data to enable application migrations, but application stacks are fully realized within a single cloud.

Cross-cloud networking introduces the concept of combining services from multiple clouds and building an application stack that has its different layers hosted in different cloud providers.



The Cross-Cloud Network effectively provides the experience of a single cloud across multiple clouds.

Google Cloud

Cross-Cloud Network provides a full stack enterprise-grade solution with low latency, strong security posture, and reliability by leveraging Google Cloud and partner products to enable distributed applications.

Cross-Cloud Network makes it easier to build and assemble distributed applications across clouds while reducing total cost of ownership by up to 40%. It does so with products such as Cross-Cloud Interconnect offering a managed interconnect with 10 Gbps or 100 Gbps bandwidth, backed with a 99.99% SLA. It supports Alibaba Cloud, Amazon Web Services, Microsoft Azure, and Oracle Cloud Infrastructure with availability in all Google Cloud regions to enable customers to drive faster business outcomes.

Google Cloud provides design guides showing how customers can deploy their distributed applications with global and regional constructs, services, and networking paradigms to ensure application performance.


Internet-facing application and content delivery

Quality of experience is crucial for internet-facing applications such as gaming, shopping, shorts, live streaming, and super apps.

Google Cloud’s global front end provides the foundation to help organizations deliver, scale, and protect their internet-facing applications using Google Cloud’s global infrastructure.

It consists of Cloud Load Balancing, Cloud Armor, and Cloud CDN which work together to protect and deliver optimal performance and user experience. The global front end is available in all Google Cloud regions and can provide a unified front end by leveraging Cross-Cloud Interconnect to bring other clouds together.

With Cross-Cloud Network, our global frontend delivers the performance, reliability, and global reach with 40% lower TCO.



There are several advantages to using a global front end including the following:

Reliability:

Improves application availability and performance with Google Cloud Load Balancing.

Performance:

Delivers low latency performance with Cloud CDN which significantly improves user experience.

Security:

Protects applications and content with ML-powered WAF via Cloud Armor which has prevented some of the largest DDoS attacks on the planet.

Lower TCO:

Reduce the TCO by discarding unwanted and malicious traffic at the edge. Leverage rate limiting in Cloud Armor to prevent costs incurred from DDoS attacks.

GFE is also programmable, enabling customers to customize logic on Cloud Load Balancing. Google Cloud introduced Service Extensions callouts which provides programmability and extensibility by letting customers use load balancers to make gRPC calls to user-managed or native cloud services during data processing. Customers can write callout extensions against Envoy's external processing gRPC API.

Service Extensions lets supported load balancers send a callout from the data processing path to extension backend services managed by the user. This helps Application Load Balancers use custom logic in the processing path.

Some common use cases for callout extensions are the following:

- Customize routing and traffic management by performing HTTP or URL redirects
- Add, remove, or modify headers or rewrite URLs based on complex application logic before forwarding traffic to the backend service
- Log custom information from payloads or custom headers to logging
- Implement custom user authentication and authorization
- Integrate security products such as API gateway security, BOT management, or WAF

Callout extensions are highly flexible and support a variety of customizations, enabling customers to perform a wide variety of functions for their environments. Google Cloud's global front end solution provides a powerful front end for application and content delivery with programmability, industry-leading security, and performance.

Hybrid Workforce

Hybrid workforce has become the norm for enterprises today with flexible remote and on-site work models. Security has become more complex, and often requires secure access service edge (SASE) solutions with security service edge (SSE). SSE solutions are being adopted by organizations to provide secure access to enterprise applications, SaaS and to help protect the distributed workforce. However, users connecting to SSE experience higher latency for private apps as SSE solutions rely on encrypted tunnels over best-effort internet links to reach private applications across clouds.

Organizations also find it difficult to bring their high-bandwidth on-premises user traffic into SSE for security inspection due to complex networking.

As a result, they often deploy firewalls on-prem instead. To help businesses standardize on a common SSE stack of their choice for securing access for all their hybrid workforce and enabling optimal user experience, Google Cloud has partnered with Palo Alto Networks with Prisma Access, Fortinet and Broadcom with Secure Web Gateway, to offer their SSE solutions natively in Google Cloud as part of Cross-Cloud Network.



35%


Businesses reduce network latency by up to 35% with Cross-Cloud Network SSE solution intergration.

Cross-Cloud Network can direct all on-prem user traffic to these SSE solutions hosted in Google Cloud. After security inspection, traffic is routed to applications in Google Cloud or over Cross-Cloud Interconnect to other clouds. Because the security stack is deployed natively in Google Cloud, there are no tunnels or overlay networks required, allowing the stack to perform at its best. As a result of the native integration of these SSE solutions into Cross-Cloud Network, businesses will gain security controls and up to a 35% reduction in network latency.

To support the secure use of both managed and unmanaged devices by the workforce, the Cross-Cloud Network integrates Chrome Enterprise Premium to provide a Secure Enterprise Browser (SEB) solution. All communications from any device are thus secured without the need to leverage agents or implement enterprise mobility management (EMM).

The Cloud Branch Office

As enterprises connect more directly to the Google Cross-Cloud Network, network services necessary for the operation of on-premises networks can be moved to the cloud. Network services such as DHCP, DNS and IPAM (DDI), perimeter firewalls, SD-WAN appliances, SSE security stacks and even portions of the IoT stack can be instantiated in the cloud to modernize the operations of the enterprise branch and leverage the flexibility and economies of scale that the cloud provides.



Summary

The Cross-Cloud Network effectively provides the experience of a single cloud across multiple clouds. It enables multiple enterprise use cases such as distributed applications and AI workloads, internet-facing content and application delivery, and a secure hybrid workforce that can be integrated into a cloud native architecture.

Please visit the Google Cloud Architecture Center, where design guides, reference architectures, and blueprints are archived, by scanning the QR code.



