Enterprise Strategy Group by TechTarget

Economic Validation

The Economic Advantage of Google Cloud's Advanced Networking Services

By Aviv Kaufmann, Principal Economic Validation Analyst

January 2023

Executive Summary

As organizations look to deliver global services by extending their private clouds into modern hybrid- or multi-cloud deployments, they must seek to understand more about the network and services offered by cloud providers. A well-designed and well-documented cloud network can provide the tools, processes, automation, and support to help make connecting, scaling, and securing applications and services less complex and more effective for network, security, and Kubernetes administrators.

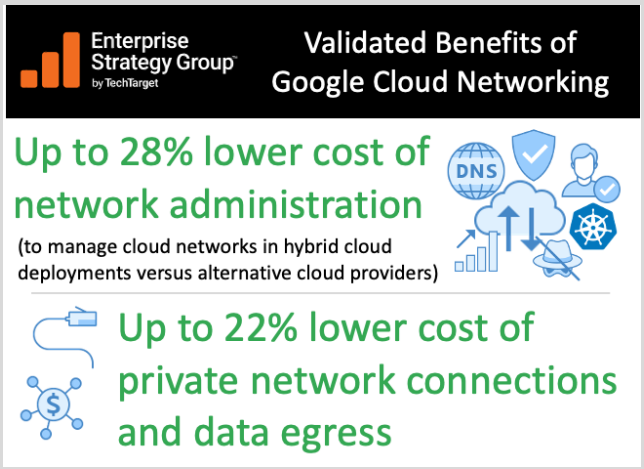

Comparing the network offerings of major cloud service providers can be confusing and is a constantly moving target but taking the time to understand the differences in the architectures and choices available can pay dividends down the line for organizations as their technology providers become their technology partners.

TechTarget's Enterprise Strategy Group validated that the Google Cloud provides a truly global network infrastructure and tools that make it simple to deploy, manage, scale, optimize, and secure hybrid cloud workloads and services. Google Cloud Networking removes much of the complexity behind cloud networks and helps to provide consistent visibility into both cloud and on-premises networks and modern Kubernetes clusters. Our models predict that Google Cloud can help lower cloud-related costs of connectivity and egress by up to 22% while lowering the expected cost of cloud network administration by 28% compared to alternative offerings.

Introduction

This Enterprise Strategy Group Economic Validation focused on the quantitative and qualitative benefits organizations can expect by using Google Cloud Networking services to connect, scale, secure, optimize, and modernize their hybrid cloud network strategy.

Challenges

As part of their digital transformation initiatives, modern organizations are moving toward a hybrid cloud strategy that allows them to extend the value of legacy applications; maintain sovereignty over their data; and take advantage of the agility, flexibility, and scale offered by public cloud services. Enterprise Strategy Group research shows that three-quarters (75%) of organizations currently use or are planning to adopt a hybrid cloud model in the next 12-24 months. But supporting applications across hybrid cloud deployments brings about challenges. Organizations may feel like they must modernize applications prior to migrating, and networking teams worry about delivering the same levels of security and performance across on-premises and cloud-based applications. Our research identifies that the top challenges or concerns with respect to supporting applications spanning public cloud infrastructure and on-premises data center infrastructure are ensuring/maintaining the proper configuration of cloud services, meeting and maintaining compliance with regulations, lack of consistent security policies, incorporating security and network configuration best practices into DevOps processes, difficulty understanding cloud security models, and a lack of visibility across different parts of the environment (see Figure 1).[1]

Figure 1. Top Ten Challenges Supporting Applications Across Hybrid Cloud Deployments

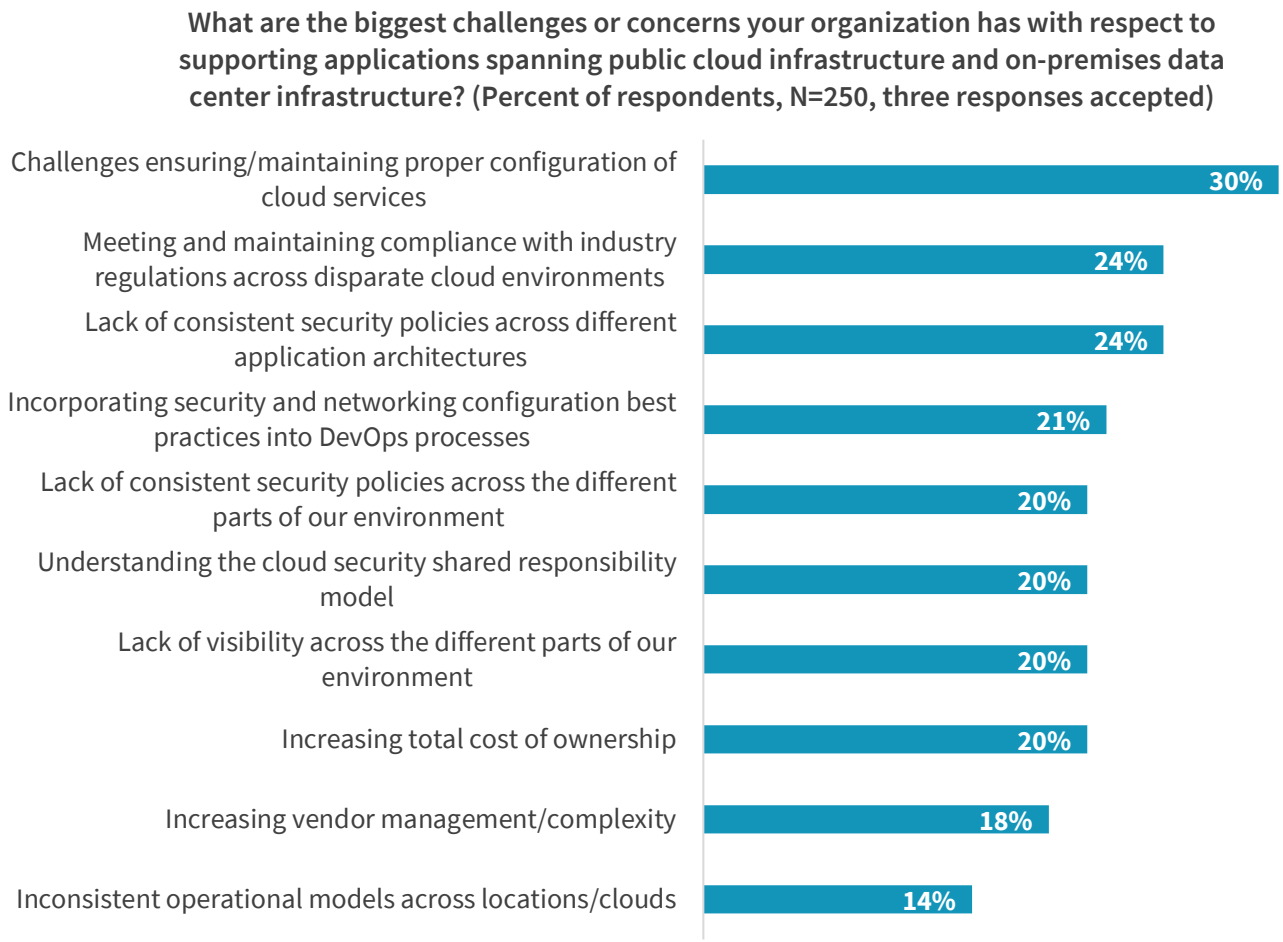

Source: Enterprise Strategy Group, a division of TechTarget, Inc.

[1] Source: Enterprise Strategy Group Research Report, *Network Security Trends in Hybrid Cloud Environments*, July 2022.

Networking in public cloud service providers is often viewed as an invisible element that simply provides connectivity. That is because early adopters are often developers or in test/dev environments where production-level requirements are not seen. Once workloads hosted by cloud service providers move into production, traditional requirements, such as service-level agreement (SLA), security, availability, consistency/predictability, observability, and simplicity of configuration and management, should be considered. Solutions to address those needs that are added after the initial design frequently lead to increased risk and may increase costs and the complexity of the architecture, as well as delay deployment. Cloud-based workloads may grow over time to encompass multiple regions to meet worldwide demands, so a simple model that worked well during a pilot phase may not scale to production.

Organizations typically struggle as the size of their public cloud environment increases over time, and with that comes changing requirements, exploding complexity, and the inevitable unintended consequences of prior decisions. It is difficult for organizations to predict how well a cloud provider's network services will meet the needs of the organization as services are delivered globally but taking the time upfront to compare and understand the differences in cloud networking offerings can pay dividends down the road.

The Solution: Google Cloud Networking Services

Google Cloud provides a networking service delivered on the world's largest dedicated, private fiber network. The same network and technologies that minimize latency and provide the end-user experience for globally consumed Google services (such as Search, Gmail, and YouTube) can be used by organizations to deliver their internal and externally facing services. This backbone currently handles 10 of the top 14 of the most popular services in the world, each of which provides services to over a billion active users. Google designed the next-generation software-defined network and distributed system technologies that are easily integrated with other Google Cloud products in a scalable, simple, and flexible manner. With over 130+ global points of presence (POPs), Google Cloud customers can leverage its low latency, dedicated global network to avoid the unpredictability of the public internet.

Google Cloud offers a broad portfolio of networking services built on top of planet-scale infrastructure that leverages automation, advanced AI, and programmability, enabling enterprises to connect, scale, secure, modernize, and optimize their infrastructure. Google Cloud Networking offers a range of benefits because it is:

Easy to Connect and Manage:

- Google Cloud Hybrid Connectivity makes it easy to connect organizations' infrastructure to the Google Cloud global network through Dedicated or Partner Interconnects, Cloud VPN, or direct and carrier peering, depending on the workload's performance and availability needs.
- **Network Connectivity Center** provides a single place to manage organizations' hybrid cloud deployments with deep visibility and tight integration with third-party solutions.
- Virtual Private Cloud (VPC) provides a single global virtual network for an entire organization that spans multiple regions without having to access the public network. Private Service Connect allows organizations to connect their private VPC network to access their own services or Google and partner service offerings.
- Cloud DNS provides reliable, resilient, and scalable DNS services that are simple to manage.

Quick to Scale:

- Cloud Load Balancing makes it easy to quickly scale applications up or down with no pre-warming needed and load balance compute resources across single or multiple regions, all while meeting high-availability requirements.

- Cloud CDN accelerates content delivery for websites and applications while lowering latency, offloading origins, and reducing serving costs. Media CDN delivers easy-to-integrate and improved streaming experiences to viewers anywhere in the world.

Highly Secure:

- **Cloud Armor** provides built-in defenses against infrastructure DDoS attacks with IP-based and geo-based access control, support for hybrid and multi-cloud deployments, preconfigured WAF rules, and Named IP Lists.
- **Cloud Firewall** provides a fully distributed, cloud-native firewall service delivering granular least-privilege control, including micro-segmentation without network re-architecture. **Cloud IDS** provides threat detection for intrusions, malware, spyware, and command-and-control attacks on your network.
- **Cloud SWG** provides a secure web gateway that helps secure egress web traffic (HTTP/S) with flexible and granular policies based on cloud-first identities and web applications.
- **Cloud NAT** allows certain resources without external IP addresses to create outbound connections to the internet while keeping your VPCs isolated and secure.
- **VPC Service Controls** provides a security perimeter for API-based services, helping to keep sensitive data private while leveraging Google Cloud storage and data processing capabilities.
- **Cloud IAP** enables secure work from untrusted networks without the use of a VPN by verifying user identity and context to guard on-premises and cloud applications.
- **Packet Mirroring** clones ingress, egress, and packet data traffic of specific instances at the VM level rather than the network level to effectively monitor and analyze security status.

Intelligent and Optimized:

- Network Intelligence Center provides comprehensive network status and monitoring of organizations' networks and the global Google network, along with proactive network verification. Centralized monitoring cuts down troubleshooting time and effort, increases network security, and allows for optimization of the overall user experience.
- Network Service Tiers allows organizations to choose to deliver prioritized traffic on the premium tier for the highest performance or cost-sensitive workloads on the standard tier for cost savings.

Hybrid Cloud Support: Google Cloud Networking versus Alternative Cloud Offerings

The purpose of this paper is to help enterprises understand the direct and indirect costs when comparing offerings between cloud service providers. While many of the costs are straightforward to compare, significant differences between cost structures and implementations may be easy to overlook. When choosing a cloud service provider to build their hybrid cloud strategies, organizations should consider the following.

1. There Are Differences between Cloud Network Architectures.

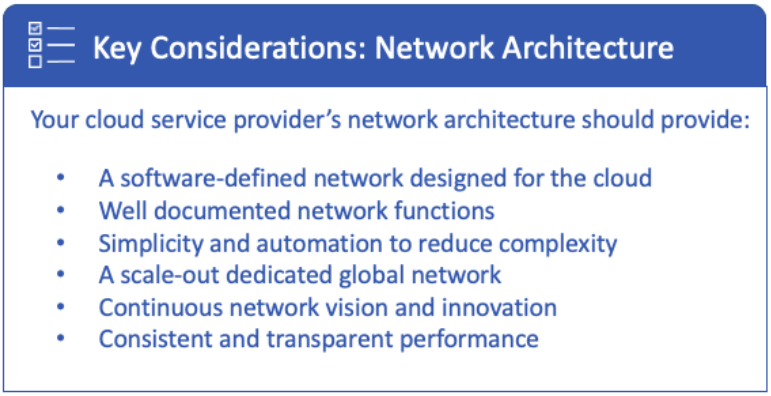

While organizations have full control over their own private networks, when employing a hybrid cloud strategy, it is important for them to understand and consider the network architecture of the cloud service provider they will be migrating applications and services to and from. While some cloud architectures may be familiar to the traditional services that on-premises network teams are accustomed to using, this may not be ideal for a modern hybrid cloud. A well-designed and well-documented software-defined network (SDN) both on-premises and in the cloud greatly reduces the administrative effort required to manage, maintain, and optimize network functions. Vendor-affiliated requests for comments (RFCs) produced for the Internet Engineering Task Force (IETF) are a good indication of the level of commitment to documenting internet standards. While almost every cloud provider offers SDN functions, there can be significant differences in how networking functions are deployed, configured, managed, and maintained. Organizations should seek to understand what functions they will need to perform, and which functions can be automated. A provider should have a well-documented SDN stack and support provisioning of a logically isolated section of the provider's cloud, with an overlay network that requires the use of SDN technology.

Additionally, organizations should make sure that their cloud service provider offers a global backbone network that helps to minimize the time that traffic is exposed to public internet connections, which results in inconsistent latency and increased risk. The network should provide predictable latency and have points of presence (POPs) in all of the locations that are important to the organization, with global load balancing and traffic that is fully encrypted and optimized for performance over the last mile. Organizations should understand any regional imitations, the difference between service offerings, egress costs, and the level of complexity that is required to configure and manage VPC policies and peering access between VPCs located in different regions.

Google Cloud provides well-documented, automated, and highly customizable SDN functions, which reduce network complexity for modernized applications and provide the customizable flexibility to handle legacy applications. Google Cloud's network is different in that it's software-defined and treated as a global feature that spans all services, simplifying network management. Google Cloud has documented its Andromeda SDN stack and is the second largest contributor to IETF RFC standards after Cisco (the largest by far of any cloud provider). When Enterprise Strategy Group compared it with other cloud providers, we found that Google Cloud's Network provided the highest network reachability of any cloud provider and reduced VPC management complexity and modeled egress costs that were roughly 21% lower than alternative cloud providers that charged for bi-directional egress between VMs. Google also provides full transparency into network latency[2] and is committed to a vision to provide continued innovation in its network and

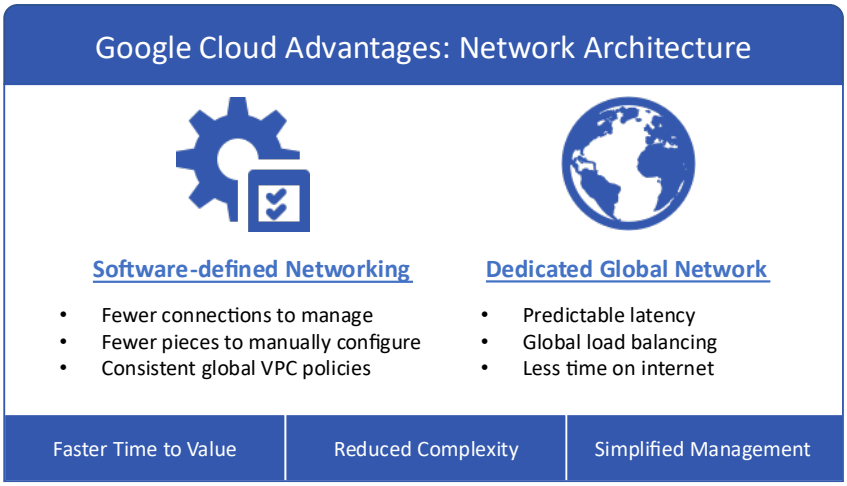

[2] *Google Cloud Inter-Region Latency and Throughput*.

data centers to deliver higher performance, efficiency, and reliability for its network services.

2. The Cloud Network Should Provide Simple and Flexible Hybrid Cloud Connectivity.

Key Considerations: Hybrid Connectivity

Your cloud service provider's network architecture should :

- Not rely solely on public internet and VPNs
- Provide private connections for performance, security, visibility, and control
- Allow use of cloud provider's global backbone as a WAN
- Limit excess management complexity and reliance on physical hardware
- Provide automation and global scope to reduce complexity

A high-performance and secure hybrid cloud network should not rely on internet connectivity and VPNs between the data center and cloud services. Most cloud providers offer several choices for connecting organizations' data centers to the public cloud. These private connections offer the highest levels of performance, security, visibility, and control. Connections can be dedicated (physical connections between the data center and the cloud provider), hosted (physical connections provided by collocated partners), or shared (hosted with shared bandwidth for each connection). The Google Cloud Network can be connected to an enterprise's multi-cloud and hybrid cloud locations using the Google Cloud Interconnect and connection points close to the organization's data center. Some public cloud providers may require the use of physical hardware gateways or that organizations lease data center space from providers that are collocated near the public cloud data center to establish direct private network connectivity. Organizations also may want to benefit from using the cloud service provider's dedicated backbone network as a WAN to connect enterprise sites located in different regions. It is important to consider the costs associated with this, as some providers charge a service fee, on top of the data transferred, to effectively do so.

Organizations should also consider the complexity involved in setting up new virtual cloud networks for projects, functions, and groups. Google is the only service provider that allows organizations to bring their own IP address (BYOIP) to any of its regions, greatly accelerating migrations, minimizing downtime, and reducing infrastructure costs. The Google network provides a global scope that spans regions and helps to automate VPC setup, subnet creation, and routing tables, providing faster time to value. Alternative cloud providers require a custom setup for each VPC and provide a regional scope rather than a global one, resulting in the need for routers and transitive gateways to manage access between regions. While subnets in any region can easily be added or isolated from a Google Cloud project, managing and documenting the configuration of VPC peering and gateways between regions for alternative solutions adds both complexity and risk to the organization. And Private Service Connect makes it easy to keep traffic secure and private on the Google network when connecting VPCs to Google Cloud, third-party, or customer-owned services without having to coordinate and perform manual tasks to negotiate and manage IP blocks and policies for each connection.

Google Cloud Advantages: Hybrid Connectivity

Flexible Connectivity Options

- Private connections, VPN
- Global scope across regions
- Fewer gateways/peerings to manage
- Easy to connect, easy to isolate

Automation and Simplicity

- Simplified migration with bring-your-own-IPs (BYOIP)
- Projects span regions
- Automated setup of services

Reduced Complexity	Simplified Management	Flexible Control

3. The Cloud Network Should Make It Easy to Scale and Accelerate Workloads.

Key Considerations: Scalability

Your cloud service provider's network should :

- Make resources easy to scale both horizontally and vertically
- Provide automated scaling, load balancing, and right-sizing of resources
- Provide intelligence and metrics that help guide scaling decisions
- Give choice in load balancing options that include health checks
- Offer high performance content delivery network to guarantee positive end-user experiences.

Once deployed and configured, the task of scaling, balancing, and accelerating workloads across organizations' hybrid cloud deployments to best meet changing demand should be simple and quick. Cloud provider resources should be easy to scale both horizontally (adding more instances of an application) and vertically (adding more resources to existing instances). Scaling through manual operations or scheduled times helps but is not as effective as automated scaling and load balancing. Automated scaling, load balancing, and right-sizing of resources make it possible to meet changing demand, ensure predictable performance, and reduce costs to the organization. Google's Compute Engine provides automated scaling of instances to predict and provide resources for individual applications as required, or to scale back instances as the workload demand is reduced. Users can set custom policies based on several metrics and use AI-generated insight to predict optimal times to scale workloads based on historical telemetry data.

Load balancing can be done at several levels, including the request and protocol level, application level, and through DNS. The ability to perform load balancing health checks is important to ensure availability at scale. Google Cloud Load Balancer (GCLB) provides server-side load balancing with SSL offload across any region and requires no prewarming. Google load balancing is simple to use and can be global, regional, or internal across private IP addresses. All traffic routing rules and policies between hundreds of services and projects can be managed centrally in a single map helping to simplify management and reduce cost for the organization.

For those businesses that are looking to speed the delivery of content across the globe, an integrated content delivery network (CDN) can be used to cache webpages, data, and media closer to consumers of the information. Besides greatly speeding up performance and avoiding bottlenecks, a CDN can provide improved reliability (more copies in more places), reduce the impact of denial of service (DoS and DDos) attacks, and help reduce costs of data transfer by reducing the amount of data that needs to be transferred. Most cloud vendors have a CDN offering, but Google Cloud also has a Media CDN for automated optimization of streaming media services.

Google Cloud Advantages: Scalability

Automated Scaling of Workloads

- Automated scaling and right-sizing based on custom metrics and predictive AI
- Simple and flexible load balancing with no pre-warming required

Accelerated Content Delivery

- CDN speeds performance and improves reliability
- Media CDN for streaming workloads

Business Agility	Operational Simplicity	Reduced Costs

These services offered by Google deliver automated scaling, load balancing, and delivery of content, which helps to greatly improve agility and reduce administrative overhead.

4. The Cloud Network Should Ensure Secure Hybrid Cloud Operations.

It is imperative for cloud providers to offer a secure platform for hybrid cloud operations, as security is a shared responsibility in a hybrid cloud model. Security is certainly top of mind for all major CSPs, and each cloud offers highly secure and available data centers, infrastructure, and technology stacks that continuously work to meet and exceed compliance standards. Organizations should insist that their provider makes use of custom security chips on every machine, effective mechanisms to establish trust between all services, and encryption in transit and at rest. As for hybrid

Key Considerations: Secure Hybrid Cloud

Your cloud service provider should provide:

- Security at the chip level and encryption in transit and at rest
- Mechanisms to establish trust between all services
- Intelligent defense against all types of attacks
- Global and regional security policies and better support for a zero-trust posture
- Visibility into the network to help identify and remediate issues faster

cloud networking, there are many vendor- and third-party-provided tools available to teams. Understanding the cost and complexity of these tools ahead of time should be a priority, as failure to do so can expose their services and even the private data center to attack. Organizations need proven and powerful tools to protect against attacks, control network access, provide support for secure access for remote workers and services, and improve visibility and control for network and security administrators.

In addition to network firewall rules, web application firewalls (WAFs) are important to providing application security and traffic filtering. Google's fully distributed Cloud Firewall can enable hierarchical policies at the global and regional levels and supports IAM-governed tags, helping to support a zero-trust posture. Google Cloud Armor was built to support hybrid and multi-cloud deployments and provides built-in ML-powered defense against infrastructure DDoS attacks with IP-based and geo-based access controls, preconfigured WAF-rules, and named IP lists. CloudArmor proved to be a critical tool for one Google customer when it was able to identify and prevent the largest layer 7 DDoS attack in history, at over 46 million requests per second.[3] Cloud IDS provides protection against and threat detection of intrusions, malware, spyware, and command-and-control attacks, with full visibility into all network traffic and VM-VM communications. Cloud NAT is a managed NAT service that helps keep Google Cloud VPCs isolated and secure, and VPC Service Controls leverage identity and context to help manage multi-tenant services and establish security parameters for API-based services. Admins can leverage packet mirroring and the Network Intelligence Center to improve visibility and speed time to identification and remediation of network-related security issues. Alternative cloud solutions may not offer these features and often require organizations to choose between several targeted options or manage and maintain multiple solutions to achieve similar protection.

Google Cloud Advantages: Secure Hybrid Cloud

Security Built for the Cloud

- Protection from DDoS attacks, intrusions, malware, spyware, and command-and-control attacks
- WAF, global firewall rules, NAT, service controls

Improved Security Visibility

- Improved visibility for security and compliance tasks
- Improved visibility into network traffic

Reduced Risk	Operational Simplicity	Improved Visibility

5. The Cloud Network Should Provide Visibility and Control.

Key Considerations: Visibility and Control

Your cloud service provider's network tools should provide:

- Unified visibility into private and public cloud networks for monitoring and troubleshooting
- Reduced complexity with a single console for network and security related services
- A visual topology of the entire hybrid cloud network
- Metrics, options, and tools that help diagnose issues and improve performance

While cloud networks are certainly designed to be simpler to manage and more automated than the traditional data center, local networking teams still greatly benefit from the ability to understand and gain visibility into and control over their cloud network and services. While the physical infrastructure running cloud services is out of their hands, it is still the responsibility of the network team to ensure the availability, security, and compliance of the entire hybrid network. Historically, visibility into some

[3] Source: Google Cloud Blog, *How Google Cloud blocked the largest Layer 7 DDoS attack at 46 million rps*, August 2022.

alternative cloud networks was limited, leading to some organizations investing in point solutions or developing custom code that helped them visualize the network topology, get access to deeper performance networks, and perform diagnostic tests. Understanding the built-in and as-a-service tools available to your team is important since it can empower your network team to avoid issues down the road and give them more flexibility.

Google's network service tiers allow organizations to optimize applications and services on the Premium Tier on a global footprint with improved performance and reliability or choose the Standard Tier, providing cost savings for workloads that do not demand the highest levels of performance. Network Intelligence Center (NIC) provides a single console solution for network observability, monitoring, and troubleshooting and supplies the insight required to help manage or reduce costs. NIC presents a visual topology of the cloud and hybrid connectivity to on-premises networks, Google-managed services, and all the associated metrics. NIC also provides monitored analysis of network configurations and firewall rules, diagnostics tools, and a global performance dashboard, helping organizations to ensure correct and consistent configuration of cloud services.

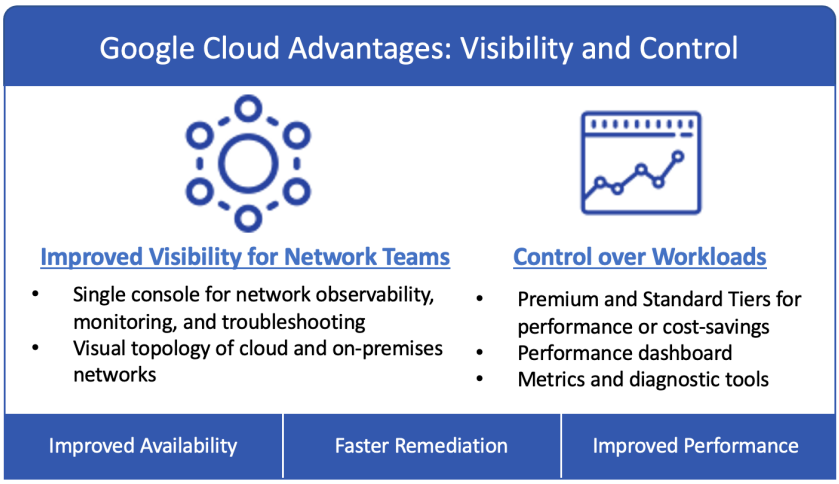

6. Initiatives Must Support Modernization Efforts.

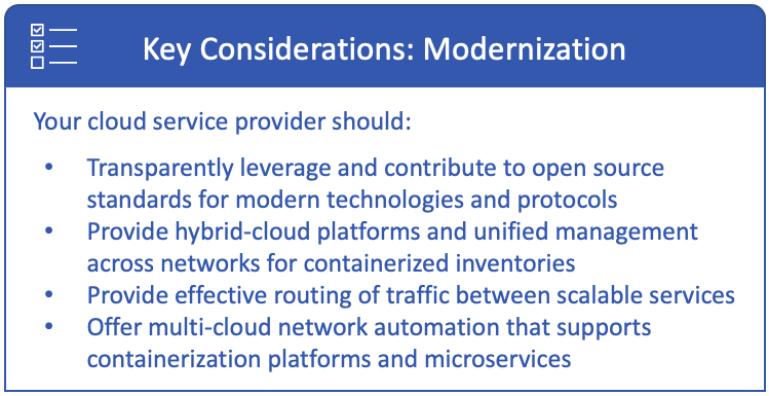

The hybrid cloud network is the backbone of modern applications. As organizations modernize their on-premises infrastructure to support containerization and begin to develop new cloud-native services or re-architect existing applications into modern and scalable microservices-based architecture, they must align strategies between on-premises and cloud networks. Cloud-based tools that provide unified management of containerized inventories between networks and infrastructure and effectively route traffic between scalable services should be a prime concern when choosing a cloud provider, even if organizations are not yet at this stage. Understanding impending needs is always important when building the foundation for the future. A cloud provider should offer network automation that supports multi-cloud containerization platforms and microservices across hybrid and multi-cloud deployments.

Google Cloud is built upon and committed to use modern open source technologies, provides transparent documentation into these technologies, and contributes back into the open source standards community. Google Cloud has been at the forefront of open source platforms, tools, and services such as Kubernetes, Envoy, Istio, knative, eBPF, web RTC, gRPC, HTTP3, and QUIC, amongst others, and will continue to innovate alongside the open source community. GKE provides a consistent networking experience between Google Cloud and on-premises Anthos deployments, as well as dynamic configuration of firewall rules, IP filtering rules, and routing tables across on-premises cluster nodes and clusters deployed in the Google Cloud. Google Traffic Director allows organizations to run microservices in a global service mesh away from their cluster by separating application and networking logic. This, coupled with DevOps practices, can result in faster time to development and greater service availability. Service Directory is a scalable managed service that provides a single place to publish, discover, and connect services, lowering the complexity of both management and operations for on-premises and cloud services.

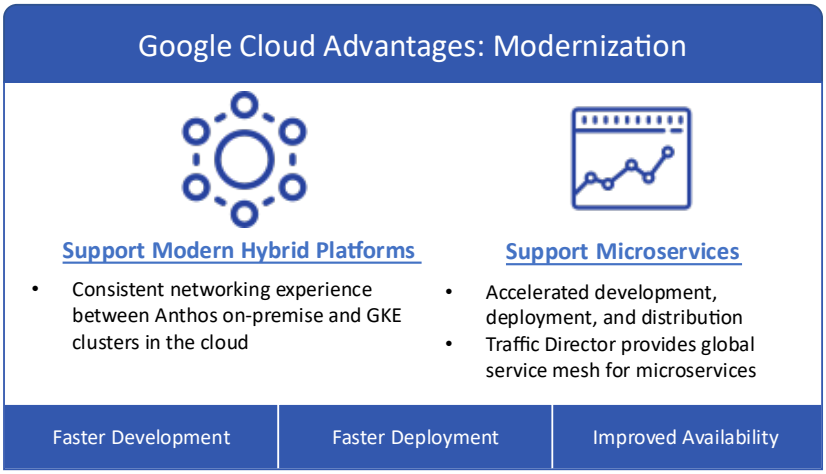

ESG Economic Validation

To validate the economic benefits of employing a hybrid cloud strategy with Google Cloud's global networking services when compared with the networks of alternative cloud providers, Enterprise Strategy Group (ESG) created requirements for a globally distributed organization and modeled some of the larger expected cloud network-related costs over a three-year period. The modeled large, globally distributed enterprise organization was looking to leverage the infrastructure within its global data center in the eastern US to deliver secure, private hybrid enterprise cloud services with predictable network SLAs.

ESG leveraged knowledge of markets, the industry, and vendor solutions as well as Google Cloud and its customer case studies to model and predict the costs to deploy, administer, manage, maintain, and operate each of the solutions. Wherever possible, direct comparisons between the solutions were used to gauge relative differences in man-hour requirements. The assumptions for the modeled analysis were carefully chosen to be as unbiased as possible so as not to give one vendor an unfair advantage, and detailed models were created to calculate and estimate the network-related costs.

Google Cloud Savings: Private Network Connectivity

We assumed that the data center was connected to the network through a redundant pair of 100Gbps private network connections. Once connected, the data center services could be made available to all instances across the global network, and cloud services could be accessed from within the data center. While connectivity through Google Cloud Interconnect and CloudRouter is simple and straightforward, the number of options and methods made available on alternative clouds can be confusing and requires that careful planning is done to determine the types of connectivity that will be needed and the physical and virtual gateway configurations that are required to best accomplish this. For this comparison, Enterprise Strategy Group simply compared the published costs of redundant 100Gbps private connections with 6 VLANs configured. The conceptual network configuration is shown in Figure 2.

Figure 2. Conceptual Network Diagram of Private Connectivity Via Dual 100Gbit links

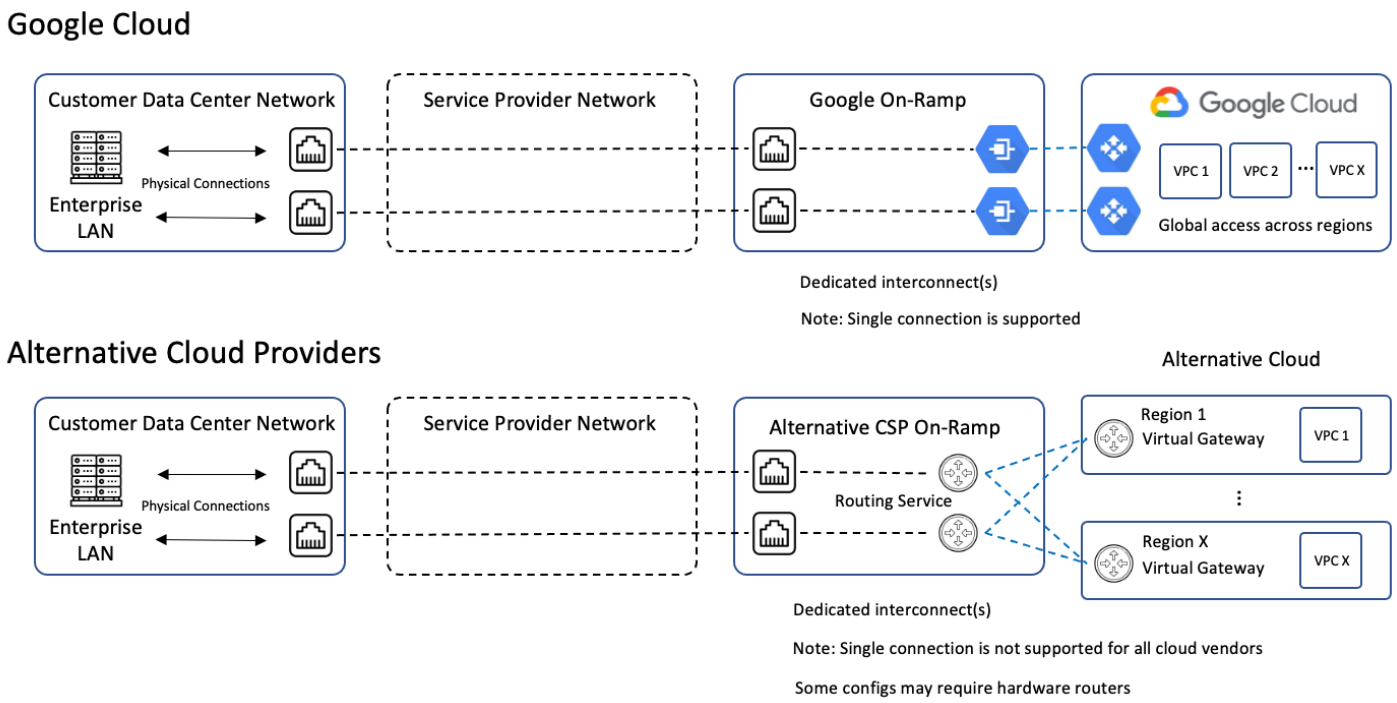

Source: Enterprise Strategy Group, a division of TechTarget, Inc.

We found that Google Cloud provided private connectivity that was 20% to 22% less expensive than the alternatives. We normalized the configurations with redundant links, but it should also be noted that Google Cloud allows for single link Interconnect for customer environments where cost savings are more critical than redundancy. Some cloud providers require the use of redundant links and in these scenarios, the customer's savings can be doubled with Google Cloud. The results of the analysis are shown in Figure 3.

Figure 3. Modeled Monthly Cost of Dual 100Gbit Direct Connections for Hybrid Cloud Connectivity

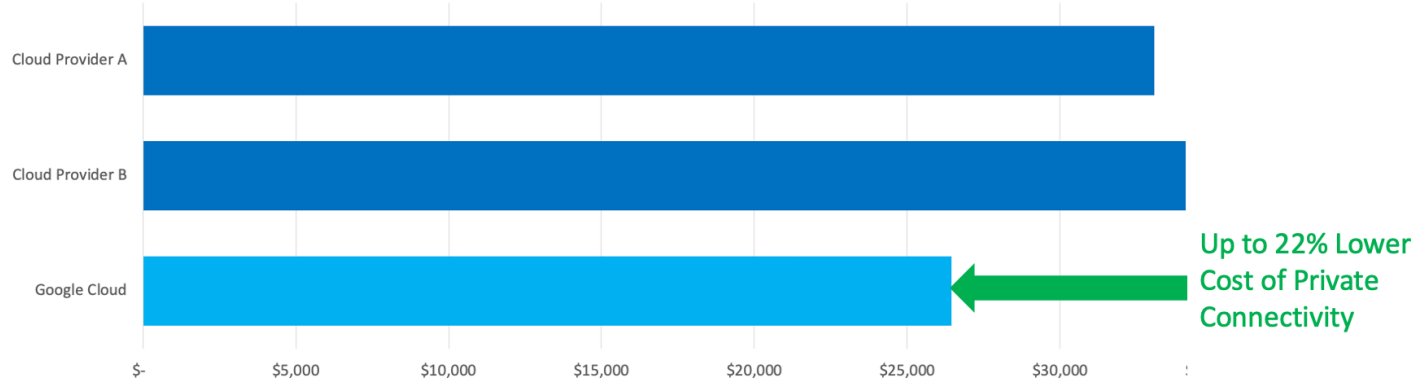

Source: Enterprise Strategy Group, a division of TechTarget, Inc.

Google Cloud Savings: Cloud Costs

Cloud costs (fees paid to the cloud provider), as might be expected from competing service providers, are generally similar in nature, but comparing all of the costs can be difficult and tedious based on tiered pricing, thresholds for free services, multiple configuration options, and uncertainty of regional locations and traffic patterns. To keep the equations as "apples to apples" as possible, we compared the expected egress rates. We assumed a total of 500TiB of traffic between instances using several types of egress. The majority would be within the same region and availability zone (200TiB) and would be

free of charge; 100TiB would be transferred to different zones in the same region; 100TiB would be transferred between regions, with half going over the global backbone (Premium Tier) and half over the internet (Standard Tier); and 100TiB would be transferred back to the data center over private connections. We found that cloud services offer similar inter- and intra-region egress rates, and similarly reduced egress rates for using private connections into their networks, but Google Cloud provides a 50% savings for inter-region transfers as they only charge for the egress from a VM and not the ingress on the other side. The results shown in Figure 4 show that Google Cloud provides up to 21% savings.

Figure 4. Comparison of Modeled Egress Charges between Cloud Service Providers

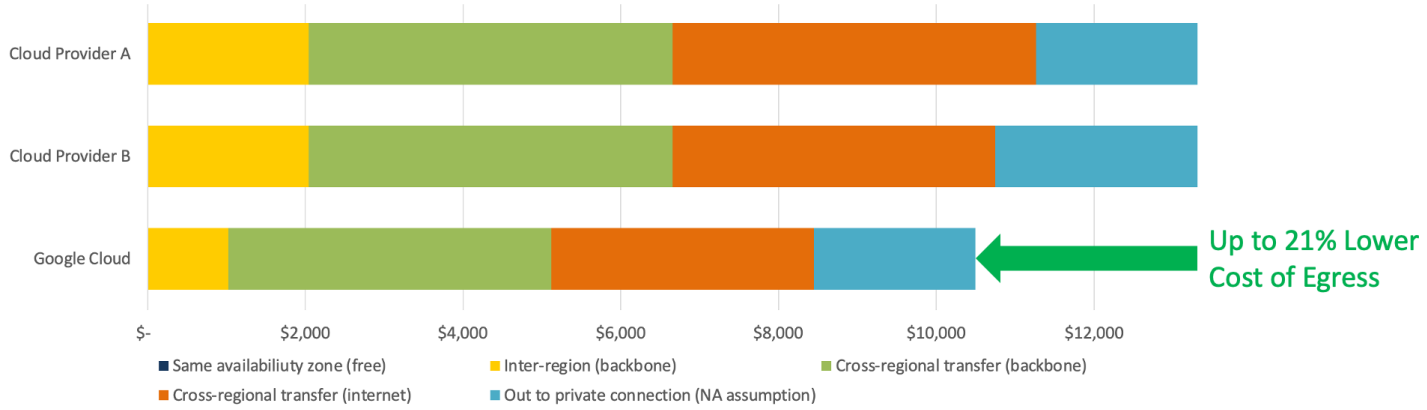

Source: Enterprise Strategy Group, a division of TechTarget, Inc.

Google Cloud Savings: Administrative Complexity

Finally, Enterprise Strategy Group (ESG) modeled and compared the costs that a large organization might expect to pay monthly to plan, purchase, deploy, operate, administer, and maintain the global network. Many of the assumptions and costs used in the scenarios were based on the reports of Google Cloud customers who have experience with multiple cloud providers' networks and could help to quantify the costs and relative differences between their deployments.

While all cloud providers offer similar tools and technologies, ESG found that Google Cloud's network was generally simpler to plan, deploy, manage, scale, secure, and monitor than the alternative solutions, as detailed in the earlier sections of this report. Detailed staff-hour models were created for each solution. The model considered the time spent by two cloud network engineers, one network security analyst, and one Kubernetes engineer to use the tools and processes provided by three cloud providers to manage, secure, and provide cloud network services for modern microservice-based hybrid cloud applications. Our models found that Google Cloud could free up 28% of the time spent by administrators, allowing them to spend more of their time planning new business requirements and testing new solutions.

Figure 5. Google Network Traffic versus AWS Network Traffic

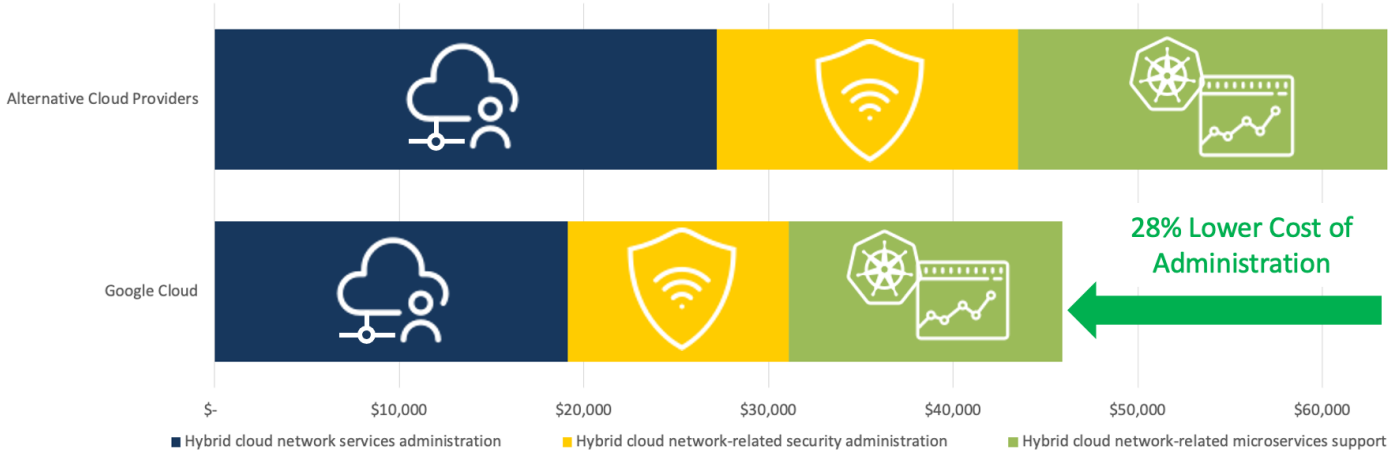

Source: Enterprise Strategy Group, a division of TechTarget, Inc.

Model Considerations

Any business consists of a set of requirements that are unique to their organization. While Enterprise Strategy Group (ESG) always makes every effort to make sure that the assumptions driving our modeled analysis are fair and conservative in nature, no single scenario could ever reflect the unique requirements of every business. ESG encourages organizations to use this study as a guide to performing their own detailed analysis based on the requirements, costs, and configurations relevant to their business.

The Bigger Truth

As organizations modernize toward connecting their data centers to hybrid and multi-cloud deployments, the relationship of the cloud service provider shifts from technology provider to technology partner. The choices made by the cloud provider related to its technology offerings, network and security design, tools, and processes can become either enablers or obstacles for network and security administrators, who are ultimately held accountable for the success or failure to support business requirements. Wise organizations consider all that they need to accomplish today and tomorrow and what a CSP can offer them before deciding on a provider.

Enterprise Strategy Group (ESG) validated that Google Cloud understands this concept and provides organizations with a platform that was designed to be truly global and simple for on-premises network admins to connect, scale, secure, and monitor. And Google Cloud offers a simple path to modernization and consistency between on-premises and cloud network functions, making it faster and easier to support both traditional workloads and modern microservices across the world.

Our models predict that Google Cloud can not only lower the cost of connectivity and egress by up to roughly 20% but can also improve the operational efficiency of your teams by up to 28% compared to alternative cloud offerings. This can free up the network, security, and Kubernetes administrators to focus on delivering new services for the business.

If your organization is looking to deliver a secure, global, high-performance enterprise hybrid or multi-cloud solution while also keeping cost and complexity to a minimum, then ESG suggests you consider Google Cloud.

Appendix: Summary of Key Considerations for Cloud Service Provider Networks

Key Considerations: Network Architecture

Your cloud service provider's network architecture should provide:

- A software-defined network designed for the cloud
- Well documented network functions
- Simplicity and automation to reduce complexity
- A scale-out dedicated global network
- Continuous network vision and innovation
- Consistent and transparent performance

Key Considerations: Hybrid Connectivity

Your cloud service provider's network architecture should :

- Not rely solely on public internet and VPNs
- Provide private connections for performance, security, visibility, and control
- Allow use of cloud provider's global backbone as a WAN
- Limit excess management complexity and reliance on physical hardware
- Provide automation and global scope to reduce complexity

Key Considerations: Scalability

Your cloud service provider's network should :

- Make resources easy to scale both horizontally and vertically
- Provide automated scaling, load balancing, and right-sizing of resources
- Provide intelligence and metrics that help guide scaling decisions
- Give choice in load balancing options that include health checks
- Offer high performance content delivery network to guarantee positive end-user experiences.

Key Considerations: Secure Hybrid Cloud

Your cloud service provider should provide:

- Security at the chip level and encryption in transit and at rest
- Mechanisms to establish trust between all services
- Intelligent defense against all types of attacks
- Global and regional security policies and better support for a zero-trust posture
- Visibility into the network to help identify and remediate issues faster

Key Considerations: Visibility and Control

Your cloud service provider's network tools should provide:

- Unified visibility into private and public cloud networks for monitoring and troubleshooting
- Reduced complexity with a single console for network and security related services
- A visual topology of the entire hybrid cloud network
- Metrics, options, and tools that help diagnose issues and improve performance

Key Considerations: Modernization

Your cloud service provider should:

- Transparently leverage and contribute to open source standards for modern technologies and protocols
- Provide hybrid-cloud platforms and unified management across networks for containerized inventories
- Provide effective routing of traffic between scalable services
- Offer multi-cloud network automation that supports containerization platforms and microservices

Enterprise Strategy Group by TechTarget

Enterprise Strategy Group is an integrated technology analysis, research, and strategy firm that provides market intelligence, actionable insight, and go-to-market content services to the global IT community.

www.esg-global.com contact@esg-global.com 508.482.0188