Demystifying the complexity of multi-cloud networking

Bhavin Desai, *Member, IEEE*, Kapil Patil, *Member, IEEE*

[1]***Abstract*— The transition of enterprise computing towards multi-cloud environments demands robust networking solutions to ensure seamless connectivity and data exchange across diverse cloud platforms. This study explores in-depth the multifaceted realm of multi-cloud networking, illuminating its complexities, motivations, and challenges faced by enterprises. By exploring the crucial role of multi-cloud strategies in optimizing workloads, enhancing cost efficiency, enabling disaster recovery, and mitigating vendor lock-in risks, we uncover the intricate balance between performance optimization and security considerations inherent in multi-cloud deployments. Moreover, the escalating adoption of artificial intelligence and machine learning (AI/ML) catalyzes a transformative shift in cloud strategies, necessitating a reassessment of network infrastructures to align with evolving demands. However, current network architectures often fall short in meeting the demands of AI/ML and multi-cloud environments, prompting the exploration of alternative solutions. To address this gap, we scrutinize three distinct approaches to multi-cloud private connectivity: Site-to-Site VPN, private connectivity through fabric providers, and direct private connectivity over fibers between cloud providers. Each approach is meticulously evaluated across dimensions such as availability, security, performance, operational maintenance, and cost-effectiveness, offering valuable insights into their merits and challenges. By illustrating the intricacies of multi-cloud networking and proposing viable solutions, this research empowers enterprises with the knowledge and tools needed to navigate the complexities of modern cloud ecosystems adeptly.**

***Index Terms*— Multi-Cloud Networking, Cloud Computing, Enterprise Computing, Site-to-Site VPN, Fabric Providers, Direct Private Connectivity, Artificial Intelligence, Machine Learning, Security, Performance, Operational Maintenance, Cost-effectiveness.**

I. Introduction

The ever-increasing adoption of multi-cloud computing is revolutionizing enterprises IT. Driven by the need to optimize workloads, reduce costs, ensure disaster recovery, and avoid vendor lock-in, a significant majority of enterprises have already implemented multi-cloud strategies. This shift is evident in diverse scenarios, from multi-cloud data lakes to SaaS applications with cross-cloud dependencies. The recent surge in AI/ML adoption is further accelerating this trend, with many organizations adding cloud vendors to leverage specialized AI/ML services. However, current network architectures often fail to meet the demands of these evolving multi-cloud and AI/ML environments. Enterprises increasingly rely on Colo-based networks, yet these alone cannot address the complex interplay of public-facing applications and cloud-specific security configurations. This paper examines the complexities of multi-cloud networking, highlighting challenges and proposing solutions that balance performance, security, operational efficiency, and cost-effectiveness to empower enterprises in navigating the modern cloud landscape.

II. Background and Shortcomings in Current Networking Architecture

The rapid adoption of multi-cloud computing is reshaping the landscape of enterprise computing. Current research indicates that 64% [1] of enterprises have already embraced multi-cloud strategies, a trend fueled by various factors, including the need for workload optimization, cost efficiency, disaster recovery, and vendor lock-in mitigation.

Some workload examples highlight the diverse motivations driving this shift. For instance, there are enterprises who operate multi-cloud data lake workloads that demonstrate a move towards a multi-cloud approach, leveraging multiple cloud service providers like Amazon Web Services and Google Cloud for their respective strengths in batch processing and data warehousing. However, while optimizing performance and cost, this introduces new challenges in inter-cloud data transfer and security.

Similarly, there are enterprises hosting SaaS applications in one cloud provider and have cross-cloud dependency, where components of a single application are distributed across multiple cloud providers. This architecture necessitates robust multi-cloud networking and security mechanisms to ensure seamless operation.

Furthermore, the recent surge in AI/ML adoption is serving as a new inflection point, compelling organizations to rethink their cloud strategies. Incorporating an additional cloud vendor to harness the capabilities of Gen AI is projected to be pursued by 31% [2] of developers. Enterprise's platform evolution illustrates

[1]Bhavin Desai is with Google, Sunnyvale, CA 94089 USA (e-mail: desai.9989@gmail.com).

Kapil Patil is with Oracle Cloud Infrastructure, Seattle, WA 98101 USA (e-mail: kapil.patil@oracle.com).

Color versions of one or more of the figures in this article are available online at http://ieeexplore.ieee.org

this trend, where the demand for multi-substrate hosting capabilities, driven by diverse end-customer requirements and the utilization of specialized AI/ML services like Generative Artificial Intelligence (Gen AI), BigQuery necessitates a multi-cloud approach.

Conventional network infrastructures inadequately align with the evolving requirements of AI/ML, and multi-cloud deployments. Enterprises are increasingly leveraging Colo-based networks to facilitate the seamless migration of their workloads from on-premises environments to the cloud. However, the reliance solely on this private network infrastructure proves to be insufficient.

The prevailing trend of constructing customized network architectures presents a conundrum characterized by operational complexity and leading to an escalated Total Cost of Ownership (TCO). Customers grapple with myriad challenges when securing traffic flows between applications and users across diverse cloud environments.

III. Solutions

A. Site-to-Site IPsec VPN:

Commencing multi-cloud connectivity through the implementation of Site-to-Site VPN presents a straightforward approach to integrating disparate infrastructural components. Enterprises can establish IPsec site to site vpn connections between cloud providers to enable multi-cloud networking with security through encryption of payloads. IPSec Site-to-Site VPN leverages IP Security (IPsec) tunnels to establish secure communication channels over the internet. Routing configurations within this framework offer flexibility, allowing for the utilization of either Border Gateway Protocol (BGP) over the IPsec tunnel or static routes to govern traffic flow. Encryption mechanisms employed within the tunnel, utilizing Advanced Encryption Standard (AES) with key lengths of either 128 bits (AES128) or 256 bits (AES256), coupled with the implementation of Diffie-Hellman groups for key exchange, ensure robust security measures.

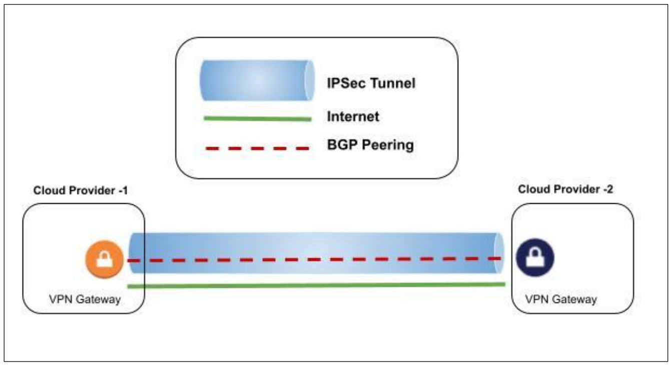

Fig. 1. Establishing Site-to-Site IPsec VPN between cloud providers with their managed VPN offering.

Important Considerations:

1) Availability and uptime:

IPsec VPNs can be configured between cloud providers using DIY approach, but the availability and uptime become the responsibility of Enterprise's operations teams and so that further involves cost. A better approach is to leverage managed offerings of cloud providers to configure Site-to-site VPN between the clouds which are backed by uptime SLA.

2) Security:

This cryptographic setup not only guarantees data confidentiality but also incorporates Perfect Forward Secrecy (PFS), augmenting the resilience of the communication infrastructure against potential cryptographic attacks

3) Performance:

IPsec VPNs are compute-intensive since it has overhead of AH and ESP protocols [3] which results in lower throughput and are limited in packets per second (pps). It also is important to consider the underlying performance of Internet connectivity between the cloud providers while measuring VPN overall performance [4]. In order to mitigate latency and enhance performance across cloud providers, enterprises can leverage multi-cloud region affinity and latency mapping to strategically deploy distributed workloads in the nearest cloud regions. By consulting performance graphs [9][10] that delineate the proximity and latency profiles of various cloud regions, enterprises can make informed decisions regarding the optimal placement of workloads, thereby minimizing network latency and maximizing overall performance across their multi-cloud infrastructure.

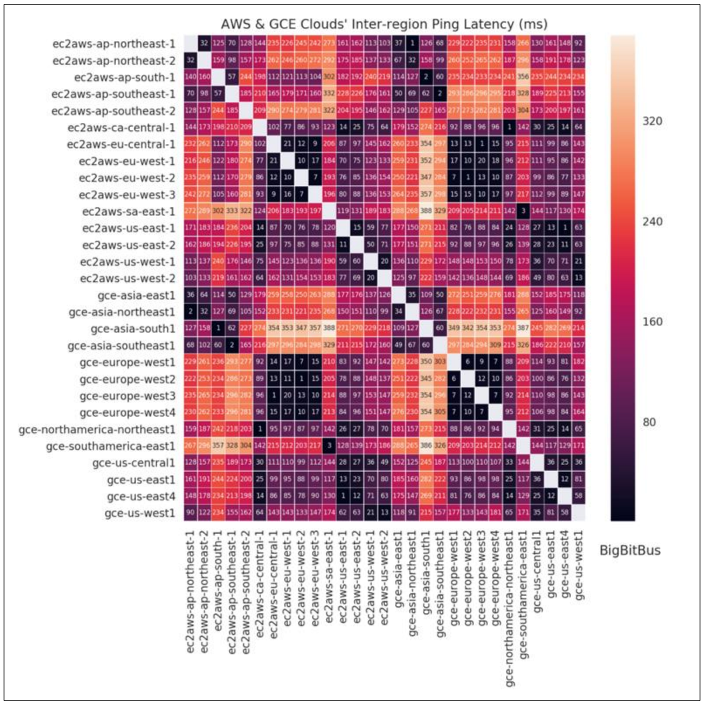

Fig. 2. Public cloud Inter-region Network latency as heat-maps [9]

AWS Regions	Azure Affinity Region Latency	GCP Affinity Region Latency
us-east-1 (Virginia)	East US 1 (Virginia), 1.87ms	us-east4 (N. Virginia), 1.91ms
	East US 2 (Virginia), 6.01ms	us-east1 (South Carolina), 14.14ms
us-east-2 (Ohio)	East US 1 (Virginia), 11.84ms	us-east4 (N. Virginia), 11.51ms
	East US 2 (Virginia), 16.52ms	us-east1 (South Carolina), 22.98ms
us-west-1 (California)	West US 1 (California), 2.17ms	us-west2 (California), 8.79ms
	West US 2 (Washington), 24.26ms	us-west1 (Oregon), 24.06ms
us-westt-2 (Oregon)	West US 1 (California), 22.93ms	us-west1 (Oregon), 13.22ms
	West US 2 (Washington), 10.78ms	us-west2 (California), 27.71ms

eu-central-1 (Frankfurt)	West Europe (Netherlands), 10.67ms	europe-west3 (Frankfurt), 1.19ms

Table. 1. Multi cloud region affinity and latency measurements by Aviatrix system [10]

4) Operational maintenance:

Very easy to setup with no CapEx or dependency on hardware. Enterprises can capture logs from cloud providers managed VPN solutions and can set up alarms based on customized metrics.

5) Cost:

IPsec Site-to-site VPNs are a very flexible and cost-effective solution to start off with. This is very easy to configure and does not require any CapEx. It also facilitates seamless decommissioning with pay as you go model in cloud providers.

B. Private connectivity through fabric providers

Several Fabric providers offers managed virtual router service tailored to facilitate direct networking interactions among workloads situated between cloud providers. In this solution, fabric providers have already laid physical fiber cross-connect in the co-location with multiple cloud providers and use their multi-tenant physical connections to carve out logical virtual connections for each customer. This solution offers enhanced flexibility and convenience in orchestrating seamless data transmission between disparate cloud environments, all without necessitating the physical presence of the Enterprise's router within the associated co-location facilities.

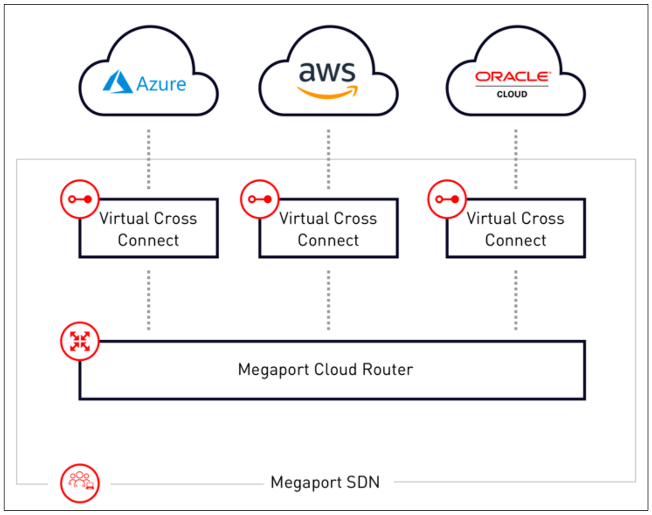

Fig. 3. Logical architecture: Multi-cloud private connectivity through Megaport – fabric provider [5]

Important Considerations:

1) Availability and uptime:

Fabric providers offer uptime SLA to the Enterprises and manage the physical and logical private connectivity offering. However, Enterprises need to consider the multi-tenant nature of the offering which can sometimes lead to oversubscription and/or noisy neighbors.

2) Security:

This cryptographic setup not only guarantees data confidentiality but also incorporates Perfect Forward Secrecy (PFS), augmenting the resilience of the communication infrastructure against potential cryptographic attacks

3) Performance:

This is a dedicated private connectivity path between the cloud providers that does not ride over the public internet. By default, the traffic is flowing in plain-text, but can be optionally encrypted using IPsec VPN. Furthermore, cloud providers also offer IEEE MAC Security 802.1AE standard (MACsec) [6] MACsec over their fiber cross-connect. MACsec establishes a robust framework for ensuring connectionless data confidentiality and integrity within media access-independent protocols. By providing comprehensive guidelines and mechanisms for safeguarding data transmission, MACsec emerges as a foundational element in fortifying network security architectures against potential threats and vulnerabilities. However, the MACsec offered is presumably still hop-by-hop, which means that unless IPsec is used on top, the customer would be trusting their plain-text data with the fabric provider.

4) Operational maintenance:

While the offering is a managed service, the deployment of virtual routers within the fabric provider network necessitates meticulous attention. This entails the negotiation and management of vendor contracts, as well as the continual maintenance of these virtual router instances. Configurations, including operating system upgrades and the establishment of Border Gateway Protocol (BGP) sessions to virtual fabric routers, demand careful orchestration to ensure seamless operation.

5) Cost:

Enterprises have to make considerations regarding capacity, bandwidth allocation, Quality of Service (QoS) configurations, and the overarching management and scalability of the network infrastructure emerge as pivotal concerns. These factors not only influence the performance and reliability of the multi-cloud private network but also significantly impact the operational expenditure (OpEx) incurred by the organization. The inherently complex nature of managing multi-cloud private connectivity, coupled with the high operational overhead necessitated by fabric provider deployments, underscores the intricate challenges encountered in this domain

C. Direct Private connectivity over fibers between cloud providers

Recent advancements in cloud networking have introduced an alternative approach to multi-cloud private connectivity, distinct from the traditional models offered by fabric providers. Leading cloud service providers, such as Google Cloud with its "Cross-Cloud Interconnect [7]" and Oracle Cloud Infrastructure (OCI) with "Oracle Interconnect for Azure [8]," have pioneered managed solutions that enable seamless private connectivity across different cloud environments without the need for intermediary hardware or routing equipment. This paradigm shift is poised to revolutionize how enterprises address the intricacies of managing their multi-cloud workloads, streamlining network connectivity and configuration processes.

In this solution, Enterprises initiate interconnects directly with their respective cloud providers and subsequently share a Letter of Authorization (LOA) with the designated managed solution provider. A physical cross-connect is then established between the routers of the participating cloud providers within a colocation facility, notably devoid of any routing hardware in the data path. This architecture ensures a private, secure, and dedicated high-bandwidth connection, meticulously provisioned and managed by the cloud providers themselves. Consequently, enterprises can reap the benefits of a simplified, high-performance multi-cloud private connectivity solution, unburdened by the complexities of network infrastructure management.

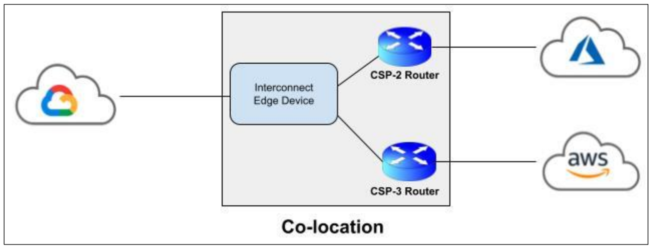

Fig. 4. Cross-Cloud Interconnect logical architecture

Important Considerations:

1) Availability and uptime:

In terms of reliability, the managed solution backs with uptime SLA to ensure service continuity and quality assurance, extending its reach seamlessly to the doorstep of other cloud providers. Moreover, its architecture promotes enhanced reliability by eliminating potential failure points, thereby fostering a robust operational environment conducive to sustained data transmission.

2) Security:

This solution's architecture is designed to minimize BGP hops, thereby enhancing network manageability and reducing potential attack surfaces. This solution supports MACsec that enables it to secure traffic between cloud providers edge routers. This streamlined approach to network management bolsters security posture and fortifies the overall resilience of the operational infrastructure.

3) Performance:

From a performance standpoint, the solution exhibits notable advantages, manifesting in reduced latency facilitated by minimizing the number of hops between Cloud Service Providers (CSPs). Enterprises can achieve line rate of the ordered dedicated bandwidth for the cross-connect and achieve predictable performance with lower latency and minimal PPS bottleneck.

4) Operational maintenance:

The solution's inherent scalability facilitates agile growth in throughput and capacity without necessitating costly upgrades to hardware in the co-location, thereby optimizing resource utilization and operational efficiency. All cloud provider metrics are published to the customers which can be used for capacity planning and operational troubleshooting as well.

5) Cost:

This solution's agility protects Enterprises from long term contractual commitments with fabric providers enabling the flexibility to adapt evolving operational dynamics and formulate robust data center exit strategies as needed. Furthermore, the abstraction of relationships with various colocation space providers streamlines procurement processes and engenders operational agility.

IV. Conclusion

In conclusion, the complexities inherent in multi-cloud networking necessitate nuanced solutions that balance security, performance, operational maintenance, and cost considerations. Site-to-Site VPNs offer a flexible and cost-effective means of initiating multi-cloud connectivity, albeit with potential performance limitations and operational overheads. Alternatively, private connectivity through fabric providers presents an enticing proposition, offering dedicated bandwidth and enhanced reliability, albeit with considerations regarding scalability and operational management. Recent advancements in cloud networking, exemplified by direct private connectivity over fibers between cloud providers, offer a paradigm shift in multi-cloud networking, promising streamlined operations, enhanced security, and optimal performance. By critically evaluating these solutions in light of their respective advantages and limitations, enterprises can navigate the complexities of multi-cloud networking more effectively, ensuring the seamless integration of diverse cloud environments into their operational frameworks.

References

[1] IDC, What Are Enterprise Multi Cloud Adoption Trends, IDC#US48902122, March 2022

[2] IDC, Developer Sentiment Indicates Generative AI is Another Driving Force for Multi Cloud Environments, Lara Greden Michele Rosen, US50858423, June 2023

[3] Performance Evaluation of VPN Protocols in Testbed. https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=a5df106a96cddfea4b02cff35f21422d23214476

[4] Namasudra S, Roy P, Balusamy B. Cloud computing: fundamentals and research issues. 2017 Second international conference on recent trends and challenges in computational models (ICRTCCM), Tindivanam, 2017, pp. 7–12. doi: 10.1109/ICRTCCM.2017.49

[5] Megaport – MCR Technical overview

[6] IEEE 802.1AE-2006 Media Access Control (MAC) Security; IEEE Standards Department: Piscataway, NJ, USA, 2006.

[7] Google Cloud blog: https://cloud.google.com/blog/products/networking/announcing-google-cloud-cross-cloud

[8] Oracle Cloud Infrastructure documentation: https://www.oracle.com/cloud/azure/interconnect/

[9] Public cloud inter-region network latency and heat map. https://medium.com/@sachinkagarwal/public-cloud-inter-region-network-latency-as-heat-maps-134e22a5ff19

[10] Multi cloud region affinity and latency. https://github.com/AviatrixSystems/Docs/blob/main/HowTos/multi_cloud_region_affinity_and_latency.rst

Bhavin Desai (Member, IEEE), As a Product Manager for Cross Cloud Network at Google, Bhavin orchestrates the journey from vision to market launch for innovative solutions and products that unlock multi-cloud magic for enterprise & strategic customers. His superpowers include product chops, go-to-market strategy, and business development magic. Bhavin has a deep expertise in cybersecurity, AI application for networking, Kubernetes & securely architecting the future.

Kapil Patil (Member, IEEE), As a Principal Technical Program Manager at Oracle Cloud Infrastructure, Kapil leads the global backbone initiatives for network capacity forecasting, planning, and scaling. With over 12+ years of experience in network engineering and cloud computing, his superpower includes architecting and deploying robust, scalable, and fortified cloud infrastructures that stand as bastions of reliability and security as well as extensive experience in deploying AI based networking and security applications cloud.