Multi-Cloud Networking

Building, securing and operationalizing enterprise cloud networks with native cloud networking and advanced network and security services

Highlights of this Market Brief:

- **The increase in enterprise cloud deployments has created demand for multi-cloud networking.** Multi-cloud network architectures must use native cloud Application Programming Interfaces (APIs) to abstract, automate, secure, and improve networking within and between multiple public clouds.
- **Multi-cloud solves problems in connecting multiple public clouds.** Enterprise organizations interviewed by Futuriom and others have expressed the need to automate intra- and inter-cloud network connectivity.
- **Growing security and compliance issues.** Managing networks presents many security and compliance challenges, all of which can be served by a well architected multi-cloud networking approach.
- **Multi-cloud networking can solve key cloud network operational challenges.** Building a common multi-cloud network data plane that leverages both native cloud networking constructs and advanced network and security services can be used to solve many operational challenges including network visibility, security, compliance, redundancy, and applications reliability**.**
- ***Multi-cloud networking use cases are growing.* Key use cases include:**
  - Multi-cloud Transit Networking
  - Network as Code Automation – Single Terraform Provider
  - Bringing Enterprise Security Policies and Next Gen Firewalls to the cloud
  - Multi-Cloud Automation and Workflow Provisioning
  - Day-two operational Visibility and troubleshooting
  - Leveraging hyperscaler infrastructure for Enterprise Virtual WAN
  - Simplified enterprise acquisition and integration
- **Abstraction-driven automation.** The future of multi-cloud networking will be driven by using APIs to connect a variety of public cloud frameworks and platforms that simplify the creation and operation of multi-cloud networks.

Intro: The Coming Multi-Cloud Networking Boom

The Information Technology (IT) industry is in the midst of its largest shift in history -- the move to cloud, powered by the need to drive digital transformation initiatives. These initiatives are transforming all organizations, ranging from governments to large business enterprises and the consumer.

The cloud has risen as the new IT platform for many reasons -- but the main reason is that it delivers a more flexible infrastructure adaptable to fast-changing demands of business applications. The enterprise transformation to cloud has become a C-Level business decision. Business leaders view moving to the flexibility and agility of cloud as the only way to fend off the existential threat of competitors who are now or will soon emerge leveraging the cloud. Well established, on-premises infrastructure and status quo application delivery is no longer a competitive edge.

Worldwide spending on public cloud services and infrastructure is expected to double over the 2019-2023 forecast period, according to the latest update to the International Data Corporation (IDC) Worldwide Semiannual Public Cloud Services Spending Guide. With a five-year compound annual growth rate (CAGR) of 22.3%, public cloud spending will grow from $229 billion in 2019 to nearly $500 billion in 2023.

There is evidence that this trend is accelerating and is still in the early stages of development. Amazon started the trend, growing revenue on its public-cloud Amazon Web Services (AWS) platform from zero in 2006 to nearly $40 billion in 2019. Now Microsoft has been catching up, reporting recent year-over-year growth rate of 62% on its Intelligent Cloud Platform, which had $10 billion in revenue in 2019. And Google is now also coming on strong, with a growth rate of 62% on its Google Cloud Platform (GCP) with revenue of $9 billion in 2019.

With all of these major public cloud vendors growing fast and increasing the capabilities of their platforms, support of multi-cloud network operations is more important than ever. How will all this data be moved, secured, and accessed by applications? It will require better multi-cloud networking. Network architects would like to simplify and automate the process of connecting to multiple clouds.

“As people need more things [from the cloud], we have become more of a multi-cloud environment," said Bobby Willoughby, a network architect with Aegon, at the recent Altitude cloud conference. "From a simplicity perspective we don’t want to do things four times. We want to do it once. Simplicity of operations is key."

Multiple clouds, or multi-clouds, will become key to IT strategy, as organizations figure out the best way to leverage and balance their applications and data in the cloudscape. But the network has fallen behind. Cloud networks will need to evolve to support the scale of data movement and processing, but the tools to connect multiple clouds are limited. Many public cloud constructs are proprietary and not concerned with connecting to clouds outside of their own service. Essentially, they are islands of networks that must be connected. As applications become increasing distributed among clouds, it will become more important to deliver automated, secure, and high-performance multi-cloud networking. IT organizations need a repeatable architecture with common services, common operational capabilities, and common security constructs across multiple cloud environments. Trying to understand and train staff to deal with the complexities of cloud networking for a single cloud is difficult. Doing this for multiple clouds is nearly impossible with the skills gap and limited resource pool. Talent is being quickly sucked up by the service providers themselves and large enterprises who have recognized the challenge and targeted local talent.

All of this means big changes for networking. Many of the common networking technologies of today were designed for closed, legacy networks built during the client-server era, prior to the arrival of public clouds. Multi-cloud networks will need to be engineered to be dynamic and automated, so that they can respond to shifting cloud workloads in real time. The network architectures, hardware centric designs, and operational models of the past -- even when extracted from hardware and deployed as virtual appliances -- are simply not viable for cloud and multi-cloud networking designs.

Let's delve into how and why this is happening.

Multi-Cloud Networking Drivers

Why will people be using multiple clouds? A wide range of survey data shows organizations want better access to multi-cloud services.

A global survey of 127 IT and business professionals, conducted by the Business Performance Innovation (BPI) Network, recently found that two-thirds of companies have now deployed enterprise applications across two or more public clouds, with 35% having moved half or more of their enterprise applications into the cloud. In addition, a recent survey by Kentik indicated that 76% of survey respondents indicated they were using cloud services, and of those respondents, nearly half (47%) were using a multi-cloud strategy.

There are a number of reasons for this, based on feedback from end users and industry trends. Here are some of the drivers:

Managing Shadow IT

Many cloud applications and development efforts start in individual business units, where cloud services are employed without going through a central IT department. This is a trend known as "Shadow IT." Shadow IT evolved to go around slow processes and draconian legacy operational approaches of the data center era. Today, the same central IT team has been put in charge of cloud operations. They recognize the past conflict of interests and want to maintain the simplicity, speed, agility, and automation that attracted Shadow IT to go around them, but they need the visibility and control they used to apply on-premises. Central IT must adopt a multi-cloud network architecture that provides the simplicity and automation of cloud and enables central IT to maintain the visibility, control, and compliance oversight they must have. A multi-cloud approach would enable various business groups to easily comply with a central IT policy while benefiting from a variety of next-generation cloud technologies.

Distributed Applications

The use of cloud applications can be quite dynamic and multi-faceted. The development group may use a range of applications in a variety of clouds to build new applications, while a specific business group might be using artificial intelligence and machine learning (ML) algorithms from a specific cloud service. Distributed applications can mean distributed clouds. Over time, this will drive the need to connect these distributed applications with smarter networks.

Some applications work better in one cloud over another other. For example:

- Enterprise Microsoft Office 365 and Active Directory on Azure
- TensorFlow AI platform on Google Cloud Platform
- Elastic Development Environments on AWS
- Oracle database applications on Oracle Cloud Infrastructure

"You are going to have workloads running in different clouds," says Amit Utreja, Sr. Prinicpal Engineer with Coupa Software. "How you expose that is a matter of how you build your networks. It's important to have that visibility."

Moving Data Closer to the Apps

As technologies such as 5G and edge computing take hold, the reach of cloud applications will expand. With applications capable of accessing data in a large variety of places, sophisticated and intelligence multi-cloud networks will be needed to access applications and data wherever they reside.

Accessing High Performance and Low Latency Data

The arrival of real-time analytics, AI, ML, and, soon, 5G means there will be more demand for Ultra Reliable Low Latency Communications (URLLC). This means that many applications -- for example, self-driving cars -- will put low-latency demands on cloud networks. Cloud providers will position resources closer to the edge of the networks and apps -- but they will still be connected to more centralized, public clouds to access data, apps, storage, and resources. The bottom line is that URLLC and edge computing will drive even more demand for multi-cloud networking connections.

Managing Against Cloud Lock-in

Many organizations don't want their IT resources and applications to be captive to one public cloud provider. However, because the various cloud providers run proprietary frameworks and are competing with one another, each cloud provider wants to keep you on their platform. IT strategists must create their own multi-cloud network architecture to plan against lock-in by driving distributed and flexible cloud strategies. A multi-cloud network architecture that creates a common data and operational control plane across multiple clouds becomes the basis for maintaining flexibility while also leveraging the best network services each cloud provides.

Redundancy and Resiliency

Most enterprise development organizations are exploring how their application workloads will be built in microservices environments that operate across clouds, offering improved security, failover, and disaster recovery. Multi-cloud networking that delivers a high-availability across multiple service providers, availability zones, and regions will facilitate the control, performance, and flexibility enterprises need in the cloud. The expansion of the cloud means that providers and enterprises will find themselves operating resources that are widely distributed geographically. Using multiple clouds can help build in redundancy into distributed applications.

Security and Compliance

Cybersecurity and IT professionals have the need for centralized visibility, control, and management of data, applications, and networks. As the center of gravity for mainstream enterprise IT moves to the cloud and away from the data center, the IT operations staff will require multi-cloud visibility and control that goes beyond what any single cloud provider can offer. They need control over security and data policies for applications running across multiple cloud platforms.

These drivers are part of a long-term trend toward distributed cloud applications and they will accelerate over the next ten years.

Multi-Cloud Networking Challenges

Some of the key questions that enterprises should ask about their multi-cloud networking platform requirements include the following:

- Multi-cloud Networking Platform: Does the multi-cloud networking platform abstract, but support, the underlying native networking constructs of major cloud providers such as Amazon, Microsoft, Google, Oracle, and others to form a common, repeatable transit network architecture that can be applied in a single or across multiple clouds?
- Common Operational Visibility: Does the multi-cloud networking platform deliver common visibility and control across multiple cloud providers?
- Security: Can the system deliver a common ingress and egress security policy for application environments (VPCs, VNETs, VCNs, etc) across clouds? End-to-end encryption inside the cloud? High-performance encryption from the data center to the cloud and inside the cloud?
- Cloud Access: Can the platform easily reconfigure existing branch office router hardware and software, without upgrade, to connect to the multi-cloud network?
- Can the multi-cloud strategy help unify networks under a single governance model?

Let's dive a little deeper into some these requirements.

The major cloud providers are happy to lend you their network and make it easier to connect into them -- but they're less than thrilled if you want to connect to competitive clouds. There are going to be challenges to building multi-cloud networks.

These challenges include:

- **Proprietary architectures.** As discussed, all the major cloud providers have their own networks, tools, and connection constructs. These are largely proprietary. Their agenda is to lock you in, not make it easy for you to architect your own enterprise multi-cloud network, and connect to a competitor's cloud.

- **Manual overhead.** If you have three or four application environments in a single cloud, manual operations are doable. But, if you have several application environments across multiple clouds automation, visibility and control is critical. While it is possible to write scripts that automate the provisioning cloud resources, day-two operations will require ongoing manual configuration of VPC / VNET routing tables, for example through proprietary cloud interfaces in order to maintain connectivity to external resources, which is time-intensive and certainly not automated.

- **Consistent security.** Inside the enterprise network, you may have consistent security standard for networks and encryption. Outside the network, most organizations rely on using standard VPNs and encryption such as IPsec. But configuring a standard security connection across clouds is not always done by default. It requires a new method of managing and automating secure and consistent multi-cloud ingress and egress security and remote access control policies across your multi-cloud network.

- **Consistent networking system.** Each cloud network is different and has different mechanisms for handling transit networking, network segmentation, load balancing, micro-segmentation, and security. Each cloud provider has a proprietary management and orchestration interface for configuring these services on their platform.

While each public cloud provider brings a rich set of common infrastructure services, the platforms are propriety in nature, which means they have their own approaches to configuration, APIs, control, and visibility. This creates a skills gap that enterprises must overcome, as each cloud platform has its own configuration, scripting, and management tools. However, these approaches can be integrated using modern multi-cloud networking platforms and techniques that create a common abstraction that simplifies multi-cloud network provisioning and ongoing operations.

Network architects have told us that they need this type of abstracted framework for managing these multi-cloud environments, to yield a consistent operational model for automating and managing multi-cloud networks.

"I would start with looking at an architecture model that can give consistency across the cloud vendors," said Luis Castillo, head of technology architecture with National Instruments. "It's important to talk about automation and focus on value. When you have automation it adds a lot of value."

A New Networking Architecture for the Cloud

The mass shift of applications from traditional enterprise and on-premises data centers to the cloud has big implications for data center and networking architectures, which will change how networks are built to serve these purposes.

There are some basic changes happening that will alter the architecture of how networks are built:

- The center of IT gravity is moving to the cloud, away from the data center. This causes an architectural shift away from traditional data center network designs to multi-cloud network designs.

- More clouds drives the need for a common network data and operational control plane so that IT has a repeatable design, and operational approach, regardless of the underlying cloud provider.

- Cloud technologies are API-driven, making the network programmable. This allows enterprises to consume cloud-native technologies and add the advanced features they need, but are not available directly from cloud providers.

- The scale of cloud connections and need for rapid reconfiguration places more emphasis on the need for a repeatable architecture, automation, and day-two operational visibility.

- Security must be built into the network data plane leveraging both cloud-native constructs (such as security groups) and advanced services (such as end-to-end and high-performance encryption for protection of data in motion).

These trends drive more complexity. For example, in order to connect to a cloud application or a Kubernetes cluster driving a globally distributed application, there is a wide range of software and equipment that needs to be orchestrated and secured in order to make it all happen right.

The big challenge is that AWS, Azure, and Google (as well as the many other public clouds) all have different proprietary architectures using their own "constructs." In order to provide a networking architecture that can "cross clouds," one needs to leverage the cloud-native functionality of each cloud, abstract that functionality with APIs, and then provide the tools to manage these connections dynamically -- or automatically. This requires a multi-cloud networking platform that can scale by leveraging the underlying native cloud networking constructs and adds advanced data plane features for high-availability, security, and day-two operational visibility and control, while at the same time natively supporting automation capabilities such as Terraform scripting.

This can be done with a series of API connections and an orchestration workflow that can manage intra- or inter-cloud connections as well as full Terraform automation.

This approach is fundamentally different from the way that client-server networks were built. In the client/server and Internet era, networks were based on enterprise switching and routing technology which used standardized protocols. Today's clouds run a large variety of proprietary, standardized, and open-source software that runs on standardized hardware. There is a larger emphasis on connecting systems using APIs which can be used to build automation into the system. Additionally, the underlying physical infrastructure that traditional data center networking teams had direct visibility and control over in the data center is completely opaque. Without advanced services, network engineers and operations teams are operating in the dark.

Cloud providers themselves provide some level of tools and frameworks that can be used by their customers. However, the objective of the cloud providers is to make their cloud sticky -- and keep the customer on their cloud as much as possible. This includes the network, which cloud providers have physically expanded massively over the years but (from an enterprise multi-cloud networking perspective) still falls short of enterprise expectations and requirements.

More traffic has moved onto the cloud-service providers networks. The cloud providers run networks that are as large as, if not larger than, the largest service provider networks. They own large core networks themselves, including transcontinental and undersea fiber capacity in the terabit range.

This isn't necessarily a bad thing (unless you are service provider). As a customer, you can leverage the cloud networks as your own. And cloud providers make it simple for customers to connect their data centers to the cloud faster. AWS Direct Connect, Azure Express Route, and Google Interconnect are well known examples.

However, connecting within and across these clouds is not always easy -- especially if you want to connect to some or all of them at once. It requires learning the specific constructs of each cloud, manual installation, and often upgrading hardware and software at branch offices.

A new approach to multi-cloud network includes using a multi-cloud network design that creates a common network data and operational control plane across multiple clouds. This type of architecture is flexible and can be adapted, regardless of the cloud providers underneath.

The diagram below, supplied by Aviatrix, shows how multi-cloud networking solutions can be used to build multi-cloud transit networks across a variety of public and private clouds, using native cloud constructs and the Aviatrix cloud network platform to deploy a multi-cloud network design.

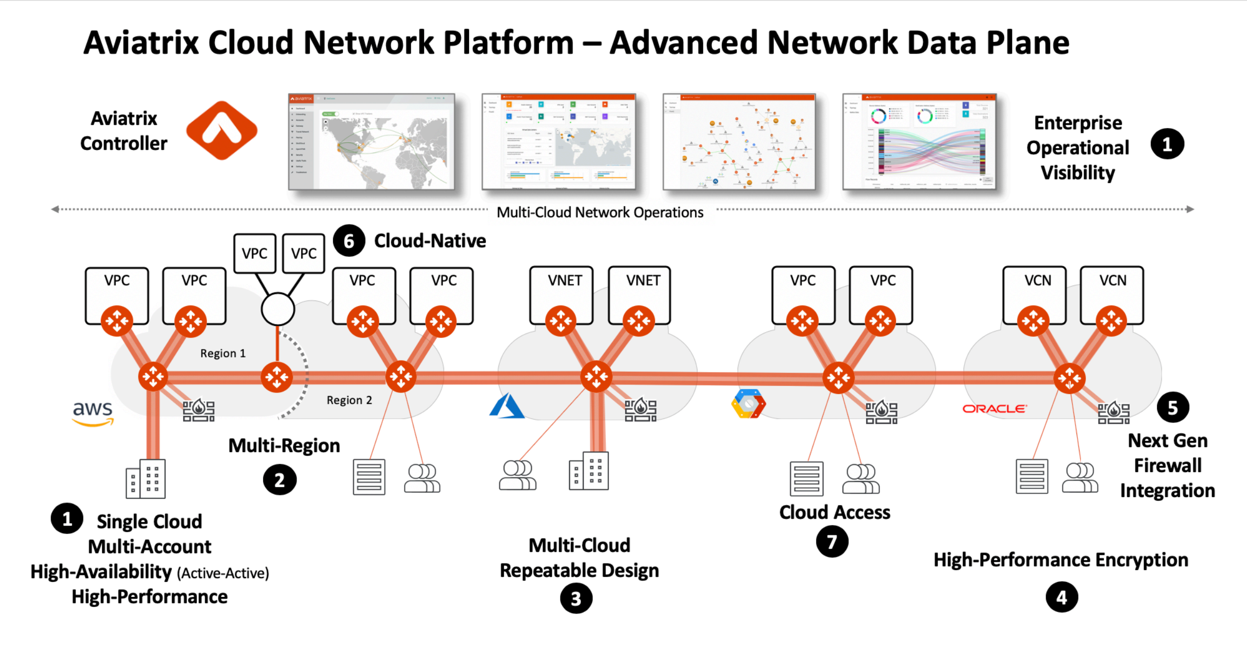
Aviatrix Cloud Network Platform – Advanced Network Data Plane
Aviatrix Controller
Enterprise Operational Visibility 1
Multi-Cloud Network Operations
6 Cloud-Native
VPC
VPC
VPC
VPC
VPC
VPC
VNET
VNET
VPC
VPC
VCN
VCN
Region 1
Region 2
aws
ORACLE
5 Next Gen Firewall Integration
Multi-Region 2
1 Single Cloud Multi-Account High-Availability (Active-Active) High-Performance
Multi-Cloud Repeatable Design 3
Cloud Access 7
High-Performance Encryption 4

Use Cases for Multi-Cloud Networking

The growing challenges of connecting application environments in multiple clouds requires a new approach. By using APIs published by the cloud providers, in combination with advance platform services available in cloud marketplaces today, it's possible to build a common abstraction layer to control, manage, and automate your own multi-cloud network design.

The best way to think about it is to look at individual use cases for the multi-cloud network, so that teams can automate repetitive tasks, gain more visibility into distributed cloud applications, and manage network connectivity inside of public cloud frameworks.

Here are some of the specific use cases that are emerging in multi-cloud networking:

Multi-cloud Transit Networks: Multi-cloud networks can be used to simplify and automate transit networking connections between application environments (VPCs, VNETs, etc.) within or between clouds. For example, this could include automating initial deployment and ongoing propagation of routes and security policies across all of your application environments. Multi-cloud networking solutions can also provide deep visibility and control over the global multi-cloud transit network, create global network segmentation policies, and automatically update route tables even for route changes on-premises that need to be propagated to VPcs/VNETs in your cloud.

Cloud Network Automation: Multi-cloud networks, as well as networks inside of a specific cloud, often require complex connection management and configuration route tables. For example, legacy data center networking requires manual configuration and management through a CLI. A modern multi-cloud networking platform, designed in the cloud for the cloud, which intelligently automates these processes through a centralized controller and APIs.

Automating Next Generation Firewalls and Security in the Cloud: Enterprise security often requires existing security policies and rules be applied in the cloud to inspect traffic from the data center to the cloud, traffic between application environments in the cloud, and Internet egress and ingress traffic. Native cloud networking constructs force compromises for deploying these traditional firewall VM appliances in the cloud. Some require IPSec VPNs, which are only used to tunnel BGP packets, not for encryption, but the IPSec processing overhead severely impacts throughput performance. Other native cloud constructs require Source Network Address Translation (SNAT), which can be used at the firewall appliance to maintain a session stickiness, but this results in the loss of source address visibility. Single firewalls may be able to overcome some of these limitations but result in other compromises, such as single point of failure or the need to pay for an expense "stand by" firewall instance that is normally an unused resource. Multi-cloud solutions can be used to automate the deployment of virtual firewall instances in the cloud, on demand; remove the need for IPSec tunneling, to maximize throughput performance; and eliminate the need for SNAT to maintain source address visibility.

Multi-cloud Workflow Networking Provisioning and Day-Two Operational Visibility: A multi-cloud network approach can be used to automate provisioning with a single Terraform provider or leverage simple workflow-based deployment processes. Leveraging the intelligence of a centralized controller and advanced network telemetry built into a multi-cloud approach allows cloud network teams to leverage day-two operational visibility across the global transit network with deep network traffic flow analysis, implement network segmentation policies, and easily troubleshoot connectivity issues between instances in the cloud, from a single console.

Connectivity Troubleshooting. Many times, network managers lack network visibility inside public clouds. An abstracted multi-cloud system can provide important network intelligence and troubleshooting tools for detecting how cloud applications and networks are interacting.

Cloud as WAN. Multi-cloud network solutions can be used to orchestrate and manage enterprise branch office routers in order to deliver a frictionless way to connect branch offices using existing router equipment with no hardware or software upgrades securely to the cloud. Once connected, enterprises can leverage the hyperscaler's backbone network to traverse the globe in the most efficient way to reach cloud-based resources or even egress to cloud-connected data centers.

User-based Cloud Access. An organization may want to build secure access to cloud applications for contractors or partners. Setting up cloud VPNs can also be a manually intensive process with compliance limitations. Multi-cloud network solutions can be used to centrally manage branch office or remote user VPN access control across clouds, simplifying the management cloud access.

These are only a handful of the growing number of uses cases for cloud networking, security, and operational control in multi-cloud networks. Use cases are growing by the month and are likely to expand substantially over time.

The Next Wave: The Abstracted Multi-Cloud

For enterprises and organizations that would like to build an integrated approach to multi-cloud networking that gives them insight, visibility, and control of networks across all cloud platforms, the answer lies in an abstracted view of the entire multi-cloud networking infrastructure.

As multi-cloud connectivity and distributed cloud applications expand, organizations will be driven toward a single multi-cloud architecture that deliver a common network design across all clouds. This architecture cannot be bound to any specific cloud provider infrastructure and must abstract the underlying cloud complexities and limitations.

The result will be centralized cloud networking control plane that lets enterprises see, control, and operationalize their cloud-based network across all of their public cloud platforms.