

SOLUTION BRIEF

Liberate Data and Modernize your Network for AI and LLM with Aviatrix

Large enterprise businesses have unique requirements for building hybrid networks. Private connectivity, on-demand capacity, high SLAs, low latency, and end-to-end encryption are at the top of the list. While these requirements have been common in regulated businesses for years, due to PII or PCI compliance needs, other verticals such as manufacturing and retail have also been adopting them as they move into global markets with strict data privacy regulations, such as GDPR. Now in the age of cloud-based AI, these requirements have become even more critical, as hybrid networks must be designed to offer the highest levels of protection, performance, and privacy.

Achieving all the network and security requirements in this arena has long been a challenge due to cost, complexity, and lack of operational visibility. Aviatrix has created the industry's first complete solution that is specifically designed for AI and LLM based workloads that require secure, high speed cloud connectivity. This solution lets enterprise businesses fulfill all the requirements that their customers and applications demand, but with but with lowered costs, greater simplicity, deep observability, and most importantly, end-to-end high performance encryption.

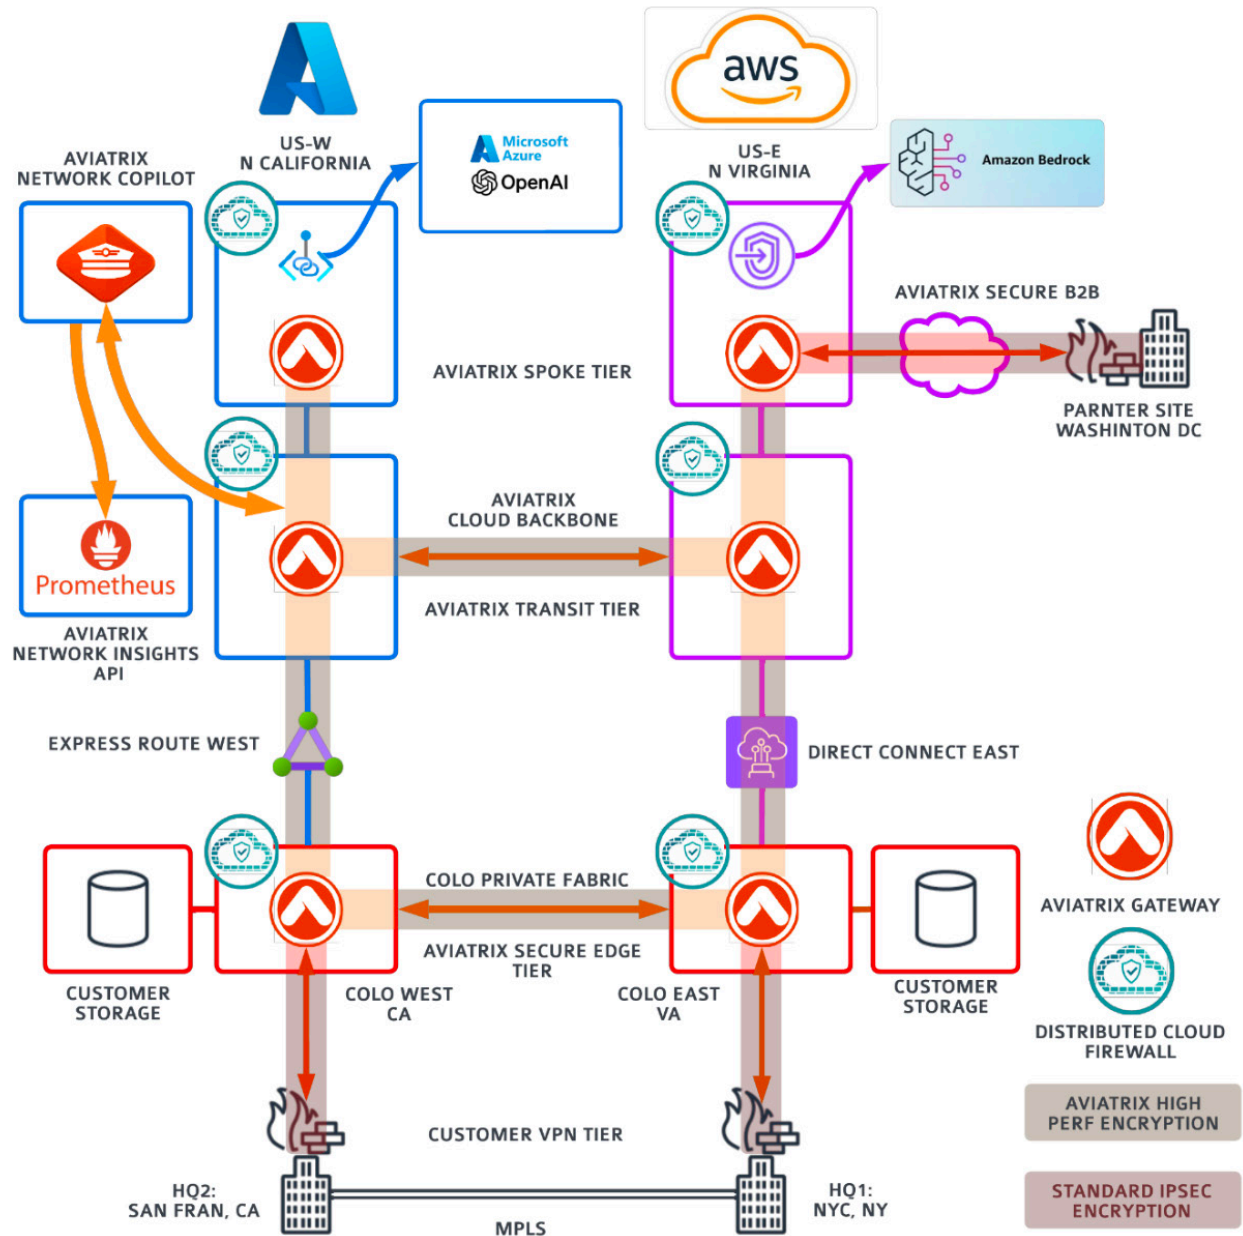
A Closer Look at Aviatrix

Before we begin unpacking how Aviatrix solves for hybrid networking for AI, ML and LLMs, let's take a quick look at how the platform works, as well as some of its most important features. This will give you a better understanding of the overall design, its benefits, and its outcomes for both your network and your business. The Aviatrix Cloud Networking Platform is a comprehensive set of integrated products that unify all cloud and customer networks together in one simple, optimized solution. The Aviatrix Cloud Networking Platform provides the features and capabilities that native cloud networks lack: agility, control, deep observability, performance, and cost savings. Aviatrix delivers these capabilities in a cloud agnostic platform that is free from provider lock in and makes customer choice and flexibility the top priority.

Aviatrix uses an ingenious combination of cloud native orchestration, automation, and built-in security to provide an innovative network platform that simplifies and optimizes both cloud-first applications and hybrid applications. The core technology that drives the platform is based on software defined networking (SDN) principles. Here a controller – the brain of the product – offloads the complexity of deploying, scaling, and securing the network, and a management console, called Aviatrix Network CoPilot, provides sophisticated telemetry, monitoring, policy enforcement, and troubleshooting tools.

Patterns for Success

To demonstrate the value of the Aviatrix solution for high performance encryption for AI and LLM, we need a framework – or a use case – to fully develop it. The use case featured here is one of many different patterns and outcomes but represents a good base line example of how many enterprises are approaching network design for connectivity to cloud-based AI and LLMs. Let's look at each tier of the design, starting at the bottom and moving up, then we will talk about the benefits that Aviatrix brings at each phase of the design.



Customer VPN Tier

The featured company, a large financial firm based in New York City, has a pair of large VPN concentrators at each building and is using them to form IPSec tunnels to a pair of Aviatrix Secure Edge gateways which sit in each colo DC. The customer has opted for two gateways in each tier, in active/active mode. The diagram shows a single Aviatrix gateway for simplicity.

The company has decided to use both clouds to help avoid lock in, improve redundancy, and provide access to two AI services to their developer teams. In this way, they can experiment with the capabilities, strengths, and limitations of each service to build the best possible service for their customers.

The customer financial data itself is not located in either HQ, but rather, is stored in a large, high performance storage array which is hosted in the colo DC by the connectivity provider. This gives the customer several distinct advantages:

1. The data can be backed up across both colo sites (east and west) for high resiliency using the private fabric of the connectivity provider. In the unlikely event of a failure to this link, the high speed, encrypted Aviatrix Cloud Backbone can be used.
2. The data is located very close to each region of importance – the west colo has a 2 ms delay to Azure US-W, and the east colo has a 3 ms delay to AWS US-E. This proximity is crucial for the LLMs in both Azure and AWS to work at optimal speeds with large quantities of data.
3. The data is housed in a remote, secure DC facility and not in the company's main office buildings. Additionally, the customer has access to a large marketplace of vendor offerings, services, and connectivity in this location.

Aviatrix Secure Edge Tier

The Aviatrix Secure Edge is a virtual machine-based solution that runs in a variety of hosted environments: VMWare, KVM, Equinix Network Edge, Megaport Virtual Edge, and within the hypervisor the Aviatrix Hardware Edge platform. Here, the company has chosen to deploy Aviatrix Secure Edge in their colo DCs to receive the following benefits:

1. **Superior performance** – Each Aviatrix Secure Edge device can handle up to 10 Gbps of encrypted traffic up to the Aviatrix gateways in the cloud, and now their cloud edge is directly adjacent to their data storage arrays in the colo, with sub-millisecond latency.
2. **Cost and efficiency** – The company has found it is far less expensive to build and operate network edge services using software as opposed to hardware. They also love the ability to scale their entire Aviatrix network on demand, including their cloud edge, and they deploy and manage it all using Terraform with a small, nimble team of two engineers.
3. **Improved security** – The company uses Aviatrix Distributed Cloud Firewall to secure their entire cloud network, and they can extend their centralized policies to their cloud edge, as the entire Aviatrix network operates like a single intelligent firewall. This gives them additional peace of mind knowing that their policy is always active across their entire network, regardless of where they build it, or how they scale it.
4. **High Resiliency and Up Time** – The Aviatrix Encrypted Fabric (or data plane) provides high levels of redundancy, especially if it is deployed with 2 or more gateways in each tier. For cloud-based gateways, Aviatrix takes advantage of Availability Zones, which are supported by redundant data centers in each cloud region. Based on this design pattern, the company has 99.99% availability, the highest for any cloud-based network, and multiple failover paths for convergence. They also enjoy the super-fast failover times offered by Aviatrix, and the fact that Aviatrix handles all gateway failover events automatically based on their HA policy.

Aviatrix Transit Tier

The company has deployed Aviatrix in a hub-and-spoke topology, which is ideal for them, as they need the hub VPCs and VNets to terminate private cloud connectivity (such as Express Route or Direct Connect), to build out centralized IT services for the cloud applications, and of course, to support multiple accounts and subscriptions that are aligned to different business units. Using Aviatrix in the Transit Tier gives them the following benefits:

1. **High Performance Encryption** – Due to PII data requirements, the customer must encrypt their customer financial data in transit from the colo all the way to their cloud AI and LLM services. By deploying Aviatrix Transit Gateways, they can extend their Aviatrix Encrypted Fabric from their Secure Cloud Edge to the cloud over their 10 Gbps private circuits at full line rate.
2. **High Speed Cloud Backbone** – As Aviatrix automatically form a high speed encrypted mesh from each gateway to another, the company quickly gained the ability to connect Azure and AWS together using the public edge of each cloud, which are directly peered in the same colos as their cloud edge. By choosing large-size VMs in each transit hub, the company enjoys 40 Gbps of encrypted bandwidth between cloud providers, which they use as a backup to their colo fabric and for redundancy between cloud apps. Cost savings on security – The company used to have a virtual DMZ using a pair of virtual firewalls in each hub but found that the Aviatrix Distributed Cloud Firewall met all their security requirements and gave them 3 times the performance at a 30% annual cost savings over their existing solution.
3. **Cost savings on network performance** – The company was also using their virtual firewalls to terminate their VPN connections from on-prem, but this presented multiple issues for them. First, they would need to double their footprint in each cloud to reach 8 Gbps of encrypted throughput, which was less than their 10 Gbps target. Second, they would also have to deploy a new VPN solution in their colos to encrypt their customer data there. When looking at Aviatrix, they found they could modernize their entire network, reach all compliance, security, and performance goals, all for a 20% annual savings on top of the 30% they saved with DCF.

Aviatrix Spoke Tier

The company has also deployed Aviatrix Spoke Gateways in each application VNet and VPC across both AWS and Azure. These VMs were right sized to provide optimal performance and cost against the 10 Gbps requirement for each AI service. By deployed Aviatrix in the Spoke Tier, the company gained the following benefits:

1. **High Performance Encryption** – As the Aviatrix Encrypted Fabric now extends all the way to the spoke application itself, the company can support 10 Gbps of encrypted through put straight into the private endpoints that connect their VNets and VPCs to their LLMs.
2. **Deep observability and cost insights** – The company uses Aviatrix Network CoPilot for visibility into their entire network, from app down to the edge. At the click of a button, their SRE team knows the health and performance of each gateway, the speed and status of each connection and tunnel, and if they want, the characteristics and meta data of every single network flow. They also use CostIQ to compute how much bandwidth each business center is using, which makes their CFO very happy.
3. **Application Data Enrichment** – The company exports all their network performance and health data from Network CoPilot to their data analytics platform using Aviatrix Insights API, where their developers can instantly correlate network behavior to their application behavior.

4. **Improved agility and flexibility** – The company has found that by deploying Aviatrix for their AI and LLM workloads, they were able to shave 2 months off their timetable for building out the supporting network and security infrastructure. After a few weeks of planning, they were able to deploy the entire end-to-end Aviatrix solution in a single day, and their IT team has almost doubled the promised amount of encrypted bandwidth to their business units.

Once their AI/LLM service in each cloud is fully trained, tested, and optimized, the company will make a careful decision about which service to use to build its new customer support application. They can then choose to keep their other AI/LLM service online for future projects or decommission it to save resources. The choice is theirs, but regardless, they will keep Aviatrix across both clouds as they know their multicloud strategy is crucial for future business decisions and outcomes.

Aviatrix Secure B2B Connectivity

Once the customer app is live, they have plans to publish it as a SaaS service to their partners. They want to use Aviatrix Secure B2B for this purpose, due to the automated connectivity, simplicity, NAT support, and built-in security. This Aviatrix solution gives them the ability to manage, operate, and troubleshoot all their partner connections from a central platform, efficiently solve for IP Overlap, and protect their network with centralized policies using Aviatrix DCF. Finally, the company does not have to make restrictions on the type of VPN device the partner provides – Aviatrix will work with **any VPN device that supports IPSec** and even support certificated-based authentication.

About Aviatrix

Aviatrix® is the cloud networking expert. We're on a mission to make cloud networking simple so companies stay agile. Trusted by more than 500 of the world's leading enterprises, our cloud networking platform creates the visibility, security, and control needed to adapt with ease and move ahead at speed. Combined with the **Aviatrix Certified Engineer (ACE) Program**, the industry's leading multicloud networking and security certification, Aviatrix empowers the cloud networking community to stay at the forefront of digital transformation. Learn more at <http://www.aviatrix.com/>.

Request a demo > <https://aviatrix.com/schedule-demo>

Get ACE Certified > <https://aviatrix.com/ace>

Learn more > www.aviatrix.com