Aviatrix Multi-Region High-Availability Cloud Network Design

Overview

As enterprises shift the center of gravity for application infrastructure from on-premise data centers to public clouds, many are designing their cloud networks on multi-region and multi-cloud architectures. Key drivers for this architectural change include:

- **Redundancy and High Availability –** Applications need Disaster Recovery (DR) capabilities. While entire region failures are rare, connectivity to a region or multiple Availability Zone (AZ) failures in same region happen. Multi-region fault tolerance and high availability are design pillars that should be part of any critical infrastructure architecture and network design.
- **Geographically Distributed Customers, Partners and Workforce –** Customers, partners and employees are physically located in multiple regions. A multi-region design allows customers, partners and employees to access cloud resources via the closest point of entry and traverse the cloud service provider's infrastructure to reach cloud-based applications and resources, wherever they reside.
- **Regional Service Availability –** Some applications are only available in certain regions of the cloud service providers.
- **Multiple Physical Data Centers –** On prem maybe in multiple regions and need closest point of entry to cloud.

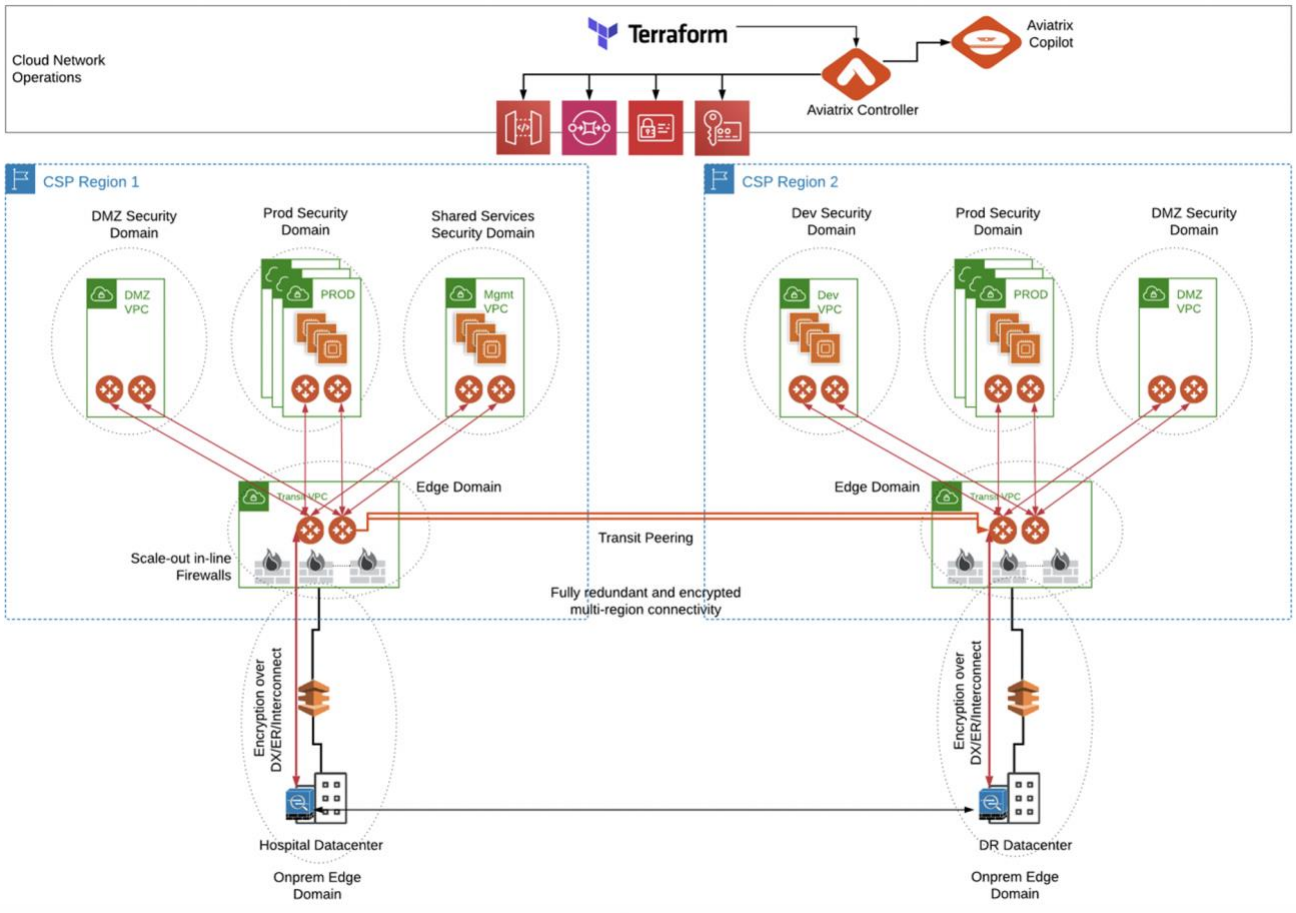

Figure 1: This Aviatrix Validated Network Design delivers a fully redundant and end-to-end encrypted, multi-region network architecture that provides high-availability, active-active connectivity failover if and when regional connectivity, availability zone or infrastructure component outages occur.

Aviatrix Introduction – Multi-Cloud Network Architecture

Aviatrix Validated Designs are created based on a Multi-Cloud Network Architecture (MCNA), which has been proven with hundreds of enterprise customers building cloud network infrastructure in AWS, Azure, Google and Oracle clouds across almost every vertical industry around the world. A MCNA is not a product, it is an architectural framework that is used as a framework for organizing design requirements. A MCNA outlines how architectural pillars such as networking, security, day-one automation and day-two operational visibility span across cloud access, cloud networking and cloud application layers. Similar to a Building architecture, the MCNA becomes the go to plan that allows cloud and security architects and their operational counter parts to work together and ensure any design meets all cross-functional requirements.

To realize their multi-cloud network designs, Aviatrix customers leverage the Aviatrix cloud network platform to deliver multi-cloud networking, security, and operational visibility capabilities that go beyond what any cloud service provider offers. Aviatrix software leverages the public cloud provider APIs to interact with and directly program native cloud networking constructs. This abstracts the unique complexities of each cloud, simplifying deployments and forming one multi-cloud network data plane with advanced networking and security features and one, consistent multi-cloud operational model. Aviatrix Transit delivers a superset of enterprise-class capabilities that becomes the foundation of our enterprise customer's multi-cloud network architecture.

Design Requirements

- Design must deliver highly available, redundant, multi-region/multi-cloud transit network
- Design must support fail-over for infrastructure component outage, AZ outage and full region outages
- Design must be flexible and utilize SDN as well as Standard network protocols and its extensions
- Design must operate in any cloud supported by Aviatrix and across clouds.
- Design must provide intelligent control to avoid overlapping CIDRs, asymmetric and route loop issues
- Design must support automated deployment using Terraform
- Design must provide enterprise-class day-two operational visibility and troubleshooting
- Design must support third-party application layer firewall service insertion

Design Elements and Features

Intelligent Centralized Controller

The Aviatrix controller is the brain of the cloud network platform. The platform leverages the centralized intelligence and knowledge of the controller to dynamically program both native cloud network constructs and Aviatrix's own gateway's advanced services. Combined with Aviatrix's Terraform provider this design enables network and security Infrastructure-as-Code automation across your multi-cloud environment.

Hub and Spoke Transit Network

Aviatrix gateways deliver advanced cloud networking and security services. Gateways are primarily deployed to deliver transit network and security services such as intelligent dynamic routing, active-active network high-availability, end-to-end and high-performance encryption and collect operational visibility data, but also for secure network ingress and egress filtering and external service insertion.

ActiveMesh – ECMP Active-Active Full Mesh Network HA

Aviatrix secure network transit is designed with active-active high-availability and redundant pathing. Pairs of Aviatrix Gateways, deployed in separate availability zones, establish a full mesh, multi-path connection that maximizes both throughput performance and network availability.

Transitive Peering

Aviatrix Transit gateways in multiple regions can be peered with each other using what is known as *Transitive Peering*. Transitive Peering leverages Aviatrix ActiveMesh to establish a full mesh, multi-path connection that maximizes both throughput performance and network availability

End-to-End and High-Performance Encryption

All traffic between Aviatrix controllers is IPSec encrypted. Standard IPSec encryption is limited to 1.25 Gbps. Aviatrix's high-performance encryption distributes processing across multiple cores and aggregates IPSec tunnels to achieve wire speed encryption, up to 75 Gbps.

Deterministic and Optimal Path Selection

ActiveMesh allows gateways to take BGP metrics such as MED, AS-PATH into consideration for the same route from different neighbors and use that to make the optimal path selection. The path selection algorithm is prescriptive and provides significant flexibility options for multi-region designs. For example, allowing for on-prem route manipulation setting preference between neighbors in different regions.

Asymmetric Routing and Loop Avoidance

To support multi-region HA, both on prem data centers advertise the same routes into multiple regions through the cloud access layer into each of the Aviatrix Transit gateways. Overlapping CIDRs heard by Aviatrix Transit gateways in both regions are advertised over the transitive peering connection to the other region's transit pair. This would normally cause serious asymmetric routing or routing loop challenges. However, the Aviatrix Controller, by default, has the intelligence to ensure the local region route is preferred over the transitive peering route, avoiding any loops or asymmetric issues.

Additional Design Considerations

- Local ASNs assigned to Aviatrix transit gateways can be modified, making it flexible and extensible.
- When using AS-PATH prepend attribute to specify traffic from on-prem data center to another region will take, be sure to account for Aviatrix Transit Gateway ASN, as Aviatrix Gateways are in the data path and will be added to route attributes.
- A MED attribute of 50 gets added on to a route when it passes through an Aviatrix Gateway, so if AS-PATH prepend does not break the tie, the MED attribute usually will.
- When leveraging Aviatrix high-performance encryption mode there are simply more tunnels between Transit Gateways. ActiveMesh views the pair a single path, next hop neighbors.
- Consider multi-region HA for other layers of your MCNA. Your application layer, for example, should be replicated across regions for DR scenarios. Your access layer should leverage multiple connections to on-prem data center with either multiple dedicated circuits, multiple internet circuits or a combination of both. Your security architecture could design in redundant firewalls in each region leveraging Aviatrix FireNet at each transit hub.

Deployment Simplicity

Multi-Region

1. Establish your Transit Network in each region using the Aviatrix Controller UI workflow or Terraform Automation.
2. Connect the Regional Transit Hubs together using the "Transitive peering" option to establish active-active, full mesh, fully redundant, encrypted multi-region connectivity.
3. Add additional features such as Transit FireNet and multi-region network segmentation as needed.

Extending to Multi-Cloud

1. Establish your Transit Network in additional cloud(s) using the existing Aviatrix Controller UI workflow or Terraform Automation.
2. Connect the Transit Hub in your first cloud with the transit hub in your second cloud together using the same "Transitive peering" option to establish active-active, full mesh, fully redundant, encrypted multi-cloud connectivity.
3. Add addition features such as Transit FireNet and multi-cloud network segmentation as needed.

Engage with Aviatrix

Online Documentation:
docs.aviatrix.com

Help with MCNA Architecture or this validated design:
info@aviatrix.com

Join the Aviatrix Community

About the Author

Hariram Sankaran is a Partner Solutions Architect at Aviatrix based in Los Angeles. Hari created this design for a customer in the Healthcare industry. This design has now been validated and adopted by dozens of Aviatrix enterprise customers.

Hari can be reached at hsankaran@aviatrix.com