AVIATRIX

DESIGN & DEPLOYMENT GUIDE

Aviatrix Edge Cloud On-Ramp Design

Best Practices for hybrid connectivity in AWS, Azure and GCP

Reference Architecture

Preface

Guide Types

Overview guides provide high-level introductions to technologies or concepts focusing on the business value.

Technology guides provide introduction to product capabilities for using Aviatrix to provide visibility, control, and protection to applications built in a specific environment. These guides describe the technologies providing examples and should be considered required reading prior to using their companion design & deployment guides.

Design & Deployment guides provide definitive guidance for different deployment scenarios, as well as procedures for combining Aviatrix technologies with third-party technologies in an integrated design. To ensure the design is reproducible, these guides provide Terraform templates for deployment of Aviatrix & third-party vendors such as CSPs. Will include best practices recommendations.

Disclaimer

The guides occasionally describe products from other companies. While the steps and screenshots were accurate at the time of publication, those companies may have since updated their user interfaces, processes, or requirements. Please refer to the external vendor website for the latest documentation.

Change History

Document Version	Aviatrix Controller Versions	Date Changed	Modified By	Change Log
0.8	7.1, 7.2 and 8.0	June 2025	Solution Architecture	Draft version
1.0	7.1, 7.2 and 8.0	July 2025	Solution Architecture	First published version

Getting the Latest Versions

Access the latest reference architecture guides at: https://www.aviatrix.com/architecture. For more information on Aviatrix EoL policies, please check the following resources:

- https://support.aviatrix.com/Aviatrix-EOL-Policy.
- https://aviatrix.com/aviatrix-end-of-sale-notice/cloudn-eos-eol-notice.

Purpose of This Guide

This guide explains how to leverage Aviatrix Edge and all its flavors to design and deploy in hybrid cloud environments. Aviatrix Edge is an enterprise grade solution that enables customers to extend the Aviatrix network to the edge for a consistent and repeatable architecture, while maintaining the management, visibility, security, and control of such environments via the Aviatrix platform.

This guide:

- Provides an overview of how Aviatrix Edge solves the challenge of building a secure and resilient architecture to facilitate Cloud On-Ramp designs.
- Links the technical design aspects of the Aviatrix Edge to the main CSPs private link connectivity best practices to achieve the highest SLA levels.
- Offers a set of decision criteria of validated deployment scenarios, along with detailed procedures for deploying Aviatrix Edge while maintaining a cohesive and integrated design.

Audience

This guide is for technical readers, including system architects and design engineers, who want to leverage Aviatrix Edge to securely connect multi-cloud environments and the enterprise data center.

This guide assumes the reader is familiar with the basic concepts of applications, networking, security, and high availability (HA). The reader should also possess a basic understanding of network and data center architectures. To be successful, you must have a working knowledge of the Aviatrix platform.

For more information on the Aviatrix platform and the installation procedures of the main Aviatrix Edge use cases, please try the following resources:

- https://aviatrix.com/ace/
- https://docs.aviatrix.com/documentation/latest/network/edge-overview.html

Introduction

This guide provides a reference design to be followed as best practices for securely connecting the enterprise data center to multi-cloud environments with Aviatrix.

The solution provides a consistent operational model to efficiently manage Aviatrix Cloud On-Ramp solution with Aviatrix Edge.

Please refer to the Glossary section at the end of this document for a full list acronyms, terms and definitions that might be referred here in this guide.

Aviatrix Edge

Aviatrix Edge enables high-performance encryption using Aviatrix patented IPsec encryption technology, providing secure, and scalable networking.

This solution is particularly beneficial for enterprises that need to manage complex, distributed networks, as it reduces operational costs, improves uptime, and accelerates cloud deployments. The following sections describe the Aviatrix Edge connectivity best practices and its differences across all four types of delivery as well as additional CSP documentation URLs for reference.

This solution guide will not present detailed step-by-step installation procedures of each one of the four delivery types of Aviatrix Edge: self-managed (ESXi or KVM), Equinix and Megaport Marketplace and AEP (Aviatrix Edge Platform) running on supported servers.

Architecture Overview

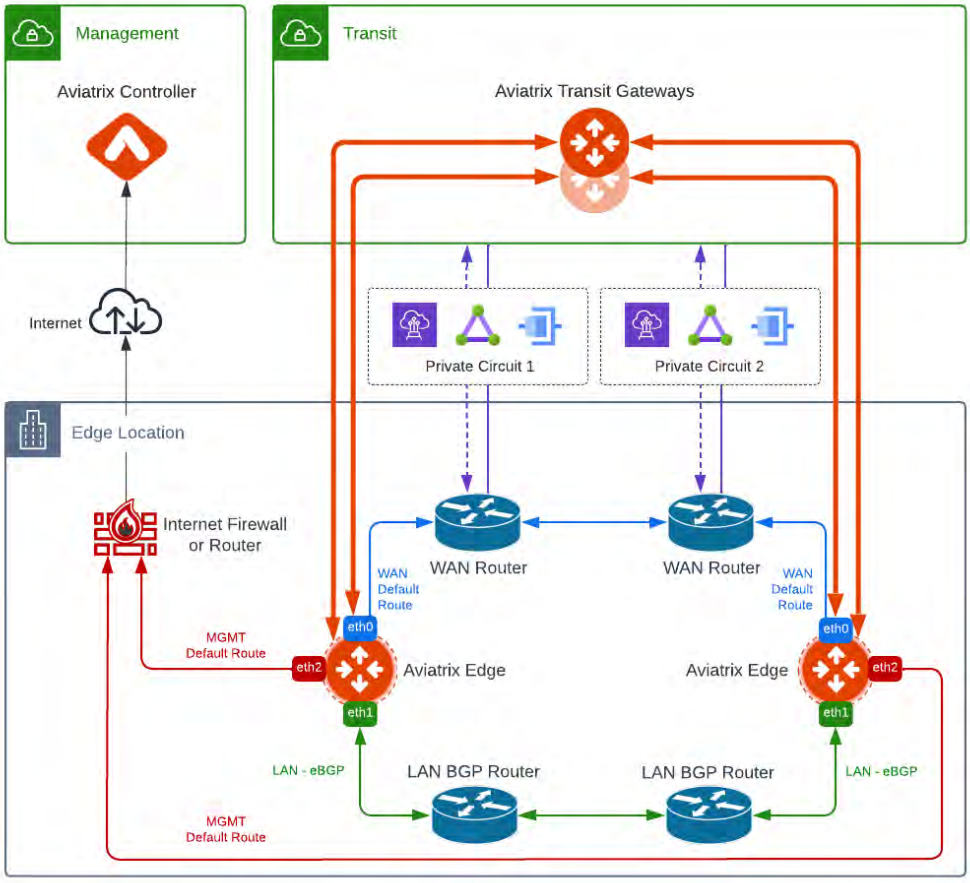

The Aviatrix Hybrid Cloud is an enterprise grade solution that that enables customers to manage hybrid cloud connections, provides deep visibility, advanced traffic engineering, and enhanced troubleshooting capabilities across hybrid cloud environments.

It also enables high-performance encryption using Aviatrix patented IPsec encryption technology, provides secure, and scalable networking and seamless connectivity to edge locations such as data centers, colocations, remote sites, and provider locations such as Equinix and Megaport.

This solution integrates seamlessly with native cloud services, such as AWS Transit Gateway, Azure Virtual WAN and GCP NCC, effectively addressing their limitations.

Aviatrix Edge Platform Options

The Aviatrix Edge solution can be delivered via different platforms to better serve each customer environment requirements and or possibilities:

- **Self-managed Platform:** it provides the flexibility to deploy Aviatrix Edge gateways as a virtual machine on self-managed x86 hardware. It currently offers VMware ESXi and KVM hypervisor support.
- **Aviatrix Edge Platform:** it is a turnkey solution that enables cloud orchestration of recommended hardware and Aviatrix Edge gateways for deployment in customer on-premises locations. For list of currently supported servers, please check our Docs page at https://docs.aviatrix.com/documentation/latest/network/edge-hardware-specs.html.
- **Equinix Network Edge:** it leverages Equinix Network Edge to deliver high performance encrypted connection to your single cloud, multi-cloud, or hybrid environments.
- **Megaport Virtual Edge:** it leverages Megaport high-speed fabric to deliver high performance encrypted connection to your single cloud, multi-cloud, or hybrid environments.

Selecting the best platform to deliver the Aviatrix Edge solution will often be guided by customer environment and their business goals. Although very similar, these platforms present some nuances on which features are supported (as of Controller version 8.0), which are highlighted in the following table:

Capabilities	Self-managed	Aviatrix Edge Platform	Equinix and Megaport
Registration	ZTP on Controller	Copilot	Copilot
Management over Private Network	Yes	No	Yes
HW Appliance vs. Virtual Form Factor	VMware ESXi & KVM	Dell	Virtual
Local Breakout for Internet	Yes	Yes	N/A
DCF (L4 only)	Yes	Yes	Yes
VRRP and VLAN	Yes	Yes	N/A
Multiple Instances per device/Hypervisor	Yes	No	N/A
ActiveMesh	Yes	Yes	Yes

BGP Routing over LAN/WAN	WAN => SDN LAN => Yes	WAN => SDN LAN => Yes	WAN BGP underlay support (AWS and Azure only) LAN => Yes
WAN Sub-Interfaces	No	No	No, but multiple WAN interfaces support as separated eth interfaces: 9x interfaces on Equinix and 4x interfaces on Megaport.
Transitive Routing (Spoke mode only)	Yes	Yes	Yes
Transit Mode	Not yet, as of Controller version 8.0	Yes	Yes

Aviatrix Edge Data Plane

Aviatrix uses standard ports and protocols for communication, overlay construction, and traffic flows. In contrast, Cloud Service Providers (CSPs) often utilize proprietary methods for traffic routing, security policies, and flow visibility.

The Aviatrix data plane uses IPsec to secure the overlay connectivity between Aviatrix gateways, ensuring encrypted communication across different network segments while Aviatrix ActiveMesh is leveraged as routing technology to maintain the connectivity between Aviatrix Edge gateways and other Aviatrix gateways via Aviatrix attachment or peerings.

For more information on ActiveMesh technology, please check our Docs page: https://docs.aviatrix.com/documentation/latest/network/activemesh-about.html.

The Aviatrix Edge will also support the following connectivity modes when peering with 3rd party devices as part of the hybrid cloud environment:

- BGP over LAN
- BGP over IPsec (Transit mode only)
- BGP over GRE (Transit mode only)
- Static routing when using VLAN interfaces (Active-Standby Spoke mode only)

For more information on how to properly setup these connections, please check our Docs page: https://docs.aviatrix.com/documentation/latest/network/edge-s2c.html.

Aviatrix Edge Control Plane

The Aviatrix Control plane is managed and maintained by the Aviatrix Controller while it communicates with all its gateways to deploy, configure, monitor, and manage the network infrastructure across various cloud environments. Such communication involves the exchange of control and management data necessary for the functioning of the data plane facilitated by Aviatrix.

The primary protocol used for communication between the Aviatrix Controller and Aviatrix gateways is HTTPS (TCP port 443), which ensures secure communication through encryption.

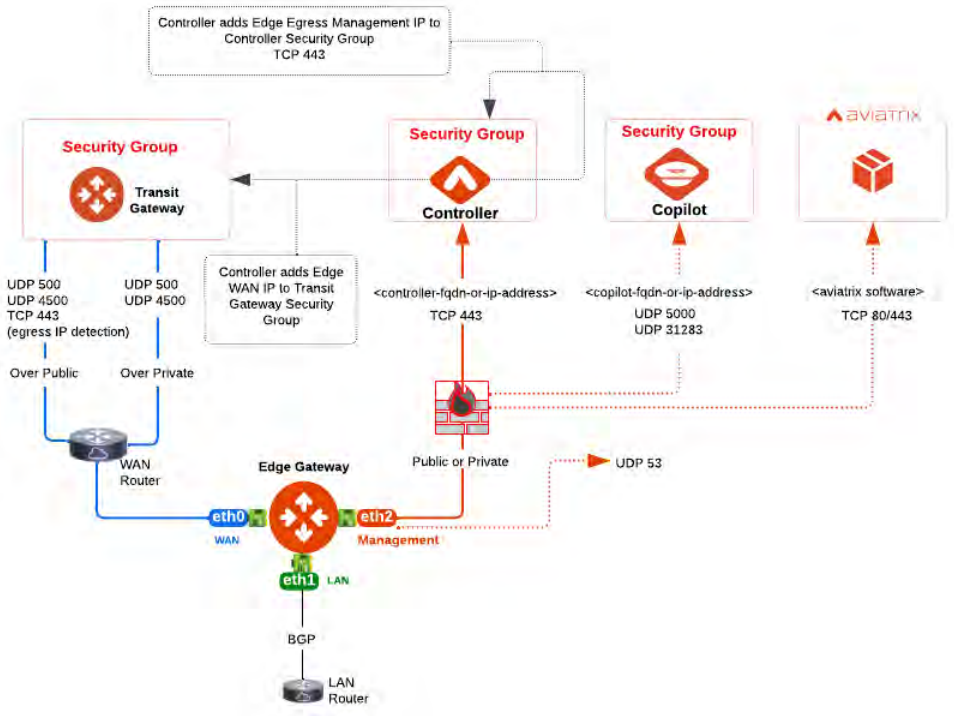

Aviatrix Edge gateways will leverage its management interface to connect to the Aviatrix Controller, CoPilot, for Aviatrix software downloads and tracelog uploads, which will then require direct connectivity via a default gateway and DNS access. Outbound-only access is required for successful registration of new Aviatrix Edge gateways into the Aviatrix Controller.

Aviatrix Edge Management Plane

Users and client access to the Aviatrix Controller and CoPilot management console is secured by HTTPS, ensuring all user interactions are encrypted and secure. Resources such as Terraform, SDKs, and other automation tools may interact with the Aviatrix Controller or CoPilot for programmatic management, configuration, and monitoring.

Terraform requests to the Aviatrix Controller are made over HTTPS. This allows Terraform configurations to securely manage and automate network resources via the Aviatrix Controller. Any external tools or SDKs integrating with the Aviatrix Controller use the REST API, secured over HTTPS. Aviatrix Edge gateways will also forward Syslog and Netflow logs to Aviatrix CoPilot, when configured to do so.

Aviatrix Edge High-Availability Modes

Aviatrix Edge gateways can be deployed in two high-available modes: **Active-Active** or **Active-Standby**.

In Active-Standby Mode, the Primary and HA Edge gateway connect to the Transit gateways with one active peering and one standby peering. Only the Primary Edge Gateway actively forwards network traffic, while VRRP VIPs will guarantee the symmetry in the path. When the Primary Edge gateway goes down, traffic is redirected to the Standby Edge gateway.

In Active-Active Mode, the Primary and HA Edge gateways connect to the Transit gateways with two active peerings. All connections established between the Edge gateways and Transit gateways perform load sharing and forward network traffic. The same is true for traffic initiated from the site/LAN side of the Edge gateways. Asymmetry paths on this mode are expected and properly handled via ActiveMesh. This is the default mode.

From a design decision perspective, the Active-Standby mode should only be preferred if the proposed On-Ramp design in question requires that the Aviatrix Edge gateway LAN peers to firewalls, so symmetric paths are a requirement or if extending on premises VLAN segments to cloud environments is desirable.

Other than that, the Active-Active option should be preferred. Another key aspect of the Active-Active mode is that it allows for horizontal scaling while traffic is forwarded from and to cloud environments leveraging ECMP via the ActiveMesh routing technology.

For more details on the specifics of each HA mode, please check this Docs page: https://docs.aviatrix.com/documentation/latest/network/edge-spoke-ha.html.

For more information on how to extend VLAN segmentation and leverage VRRP to maintain Active-Standby high-availability, please check this Docs page: https://docs.aviatrix.com/documentation/latest/network/edge-spoke-vlan-segmentation.html.

From a routing integration aspect, in which the On-Ramp design to which the Aviatrix Edge is deployed, we will in the next sections of this guide explore the best practices of each CSP and then follow up with the corresponding Aviatrix Edge setup that complies with the same goals and targets. This means that we should favor the following overall conditions:

- Active-Active mode with scaled out gateways deployed in different gateways groups.
- The horizontal scale size will be directed by the contracted bandwidth of the gateways and the current benchmark of the platform being considered. For example: a pair of geo-redundant AWS Direct Connect links totaling 50Gbps would require 2x sets of 2x AEP Edge gateways running on Dell R450 servers, where the total number of AEP devices equals the total desired capacity divided by the aggregated benchmark of a single device, multiplied by the number of sites or locations considered.
- LAN connectivity should favor resiliency by establishing cross-BGP sessions with the site adjacent L3 switch stack (VRRP/HSRP). If a redundancy greater of N+1 is considered on the switch stack, the Aviatrix should be scaled at the same rate to honor the same redundancy targets and prevent bottle necks or single point of failures.
- Underlay BFD to improve resiliency and recovery from private link failures, whether handled by on premises switches routers or by the Edge gateways themselves. It is also desirable to properly configure the L2/L3 underlay to prevent unnecessary failovers on the overlay.

- Finally, proper BGP controls such as local preference can be used in the underlay to better distribute the traffic on the contracted Active-Active private links.

Aviatrix Edge Horizontal Scaling

Horizontal Scaling requires all gateways to be configured in Active-Active mode, to be part of the same Site ID and to use the same ASN. Although, each gateway needs to be part of a different gateway group. This setup allows all Edge peers to access all its routing paths via ActiveMesh technology, whether Aviatrix attachments or the BGP sessions with local peers are being considered.

As a design decision, horizontal scaling of the Edge gateways will improve the resiliency and increase the allocated bandwidth capacity associated with a single gateway instance, whether it is a virtualized self-managed node, or running in Equinix/Megaport, or in supported AEP hardware (whose capacity will be tied to the physical SFP ports being considered).

For more information on how to setup horizontal scaling, please check our Docs page: https://docs.aviatrix.com/documentation/latest/network/edge-spoke-scaling.html.

AWS, Azure and GCP Best Practices Overview

AWS Direct Connect Best Practices for Max SLA Targets

AWS recommends connecting from multiple data centers for physical location redundancy, also leveraging redundant hardware and telecommunications providers.

Additionally, AWS recommends the usage of dynamically routed, active/active connections for automatic load balancing and failover across redundant network connections, as well as provisioning sufficient network capacity to ensure that the failure of one network connection does not overwhelm and degrade redundant connections.

Multi-Site Direct Connect Design

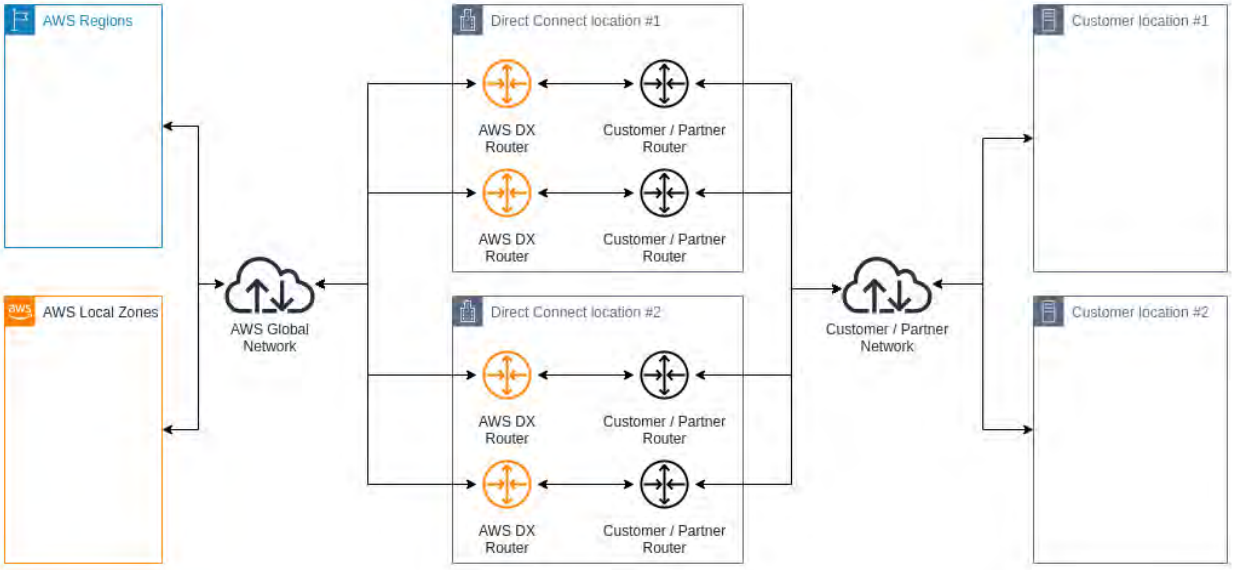

AWS highlights in its best practices recommendations that maximum resilience is achieved by separate connections terminating on separate devices in more than one location, as shown in the picture above, whose topology provides resilience to device failure, connectivity failure, and complete location failure.

More details on the SLA targets of such design can be found at the AWS specific SLA page: https://aws.amazon.com/directconnect/sla.

From an Aviatrix perspective, the multi-site best practice can be met by the following design decisions on how to contract and deploy AWS Direct Connect:

- At least a pair of Active-Active dedicated or hosted Direct Connect private VIFs per location.
- All Direct Connect VIFs should be associated with a Direct Connect Gateway, so the underlay can be extended to up to 30 VPCs. This would correspond to the number of Transit VPCs managed by Aviatrix globally considering the same Direct Connect Gateway.
  For more information on such quotas, please check the AWS specific quota page: https://docs.aws.amazon.com/directconnect/latest/UserGuide/limits.html.
- This setup allows for each site location to connect directly to any regions in AWS, as long such VPCs have their attached VGWs associated with the DXGW (Direct Connect Gateway in question).
- The AWS feature SiteLink could still be explored in case branch-to-branch connectivity is desired via Direct Connect links and could also be leveraged by Edge gateways in Transit mode using Aviatrix Transit peerings.
  For more information on AWS SiteLink, please check the following page: https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-aws-direct-connect-sitelink.

AWS BGP Settings, BFD and DXGW

While leveraging Aviatrix Edge overlay connectivity as part of the hybrid cloud connectivity, the underlay connectivity over AWS Direct Connect needs to provide the basic connectivity between the main interface of the Aviatrix Transit gateways in the target VPC (whose attached VGW is associated with the DXGW in question) and the on prem router granting the Aviatrix Edge gateway access to such environment.

This basically means two things:

- The DXGW will have to advertise only the CIDR prefixes that include the Transit gateways which the Edge gateways need to attach to. The Direct Connect Gateway quota currently supports up to 200 prefixes, which is more than enough assuming it can extend connectivity to up to 30 VPCs.
- The BGP peer on the customer side will have to advertise at minimum the WAN subnets hosting all the Aviatrix Edge gateways that are supposed to attach to the aforementioned Aviatrix Transit gateways.

From the perspective of load balancing traffic, we recommend the usage of either manual BGP local preference per group of Edge gateways or the pre-defined AWS BGP communities for the same purpose, so the underlay is better allocated by the overlay sessions.

AWS will load balance traffic towards on premises prefixes using ECMP across the active connections regardless of their home Region associations. Also, it is important to note that local preference BGP community tags are evaluated by AWS before any AS_PATH attribute and are evaluated in order from lowest to highest preference (where highest preference is preferred).

The following local preference BGP community tags are supported by AWS:

- 7224:7100 – Low preference
- 7224:7200 – Medium preference
- 7224:7300 – High preference

For more information on AWS BGP settings support Direct Connect, please check the following page: https://docs.aws.amazon.com/directconnect/latest/UserGuide/routing-and-bgp.html.

Finally, AWS recommends enabling BFD for fast failure detection and failover when connecting to AWS over Direct Connect connections. Enabling BFD for your Direct Connect connection allows the BGP neighbor relationship to be quickly torn down. Otherwise, by default, BGP waits for three keep-alives to fail at a hold-down time of 90 seconds.

Asynchronous BFD is automatically enabled for Direct Connect virtual interfaces on the AWS side. However, the on premises or COLO router needs to be configured accordingly.

For more information on how to configure BFD on most router vendors, please check the following AWS page: https://repost.aws/knowledge-center/enable-bfd-direct-connect.

For more information on how to configure BFD on Aviatrix Edge, in scenarios where Edge is also the underlay peer, please check the following Docs page: https://docs.aviatrix.com/documentation/latest/network/bgp-bfd.html#enable-bfd.

Azure ExpressRoute Best Practices for Max SLA Targets

Azure recommends the same approach to maximize resiliency of an ExpressRoute deployment to withstand failures and quickly recover from disruptions, while taking advantage of the better guaranteed SLA.

Multi-Site ExpressRoute Design

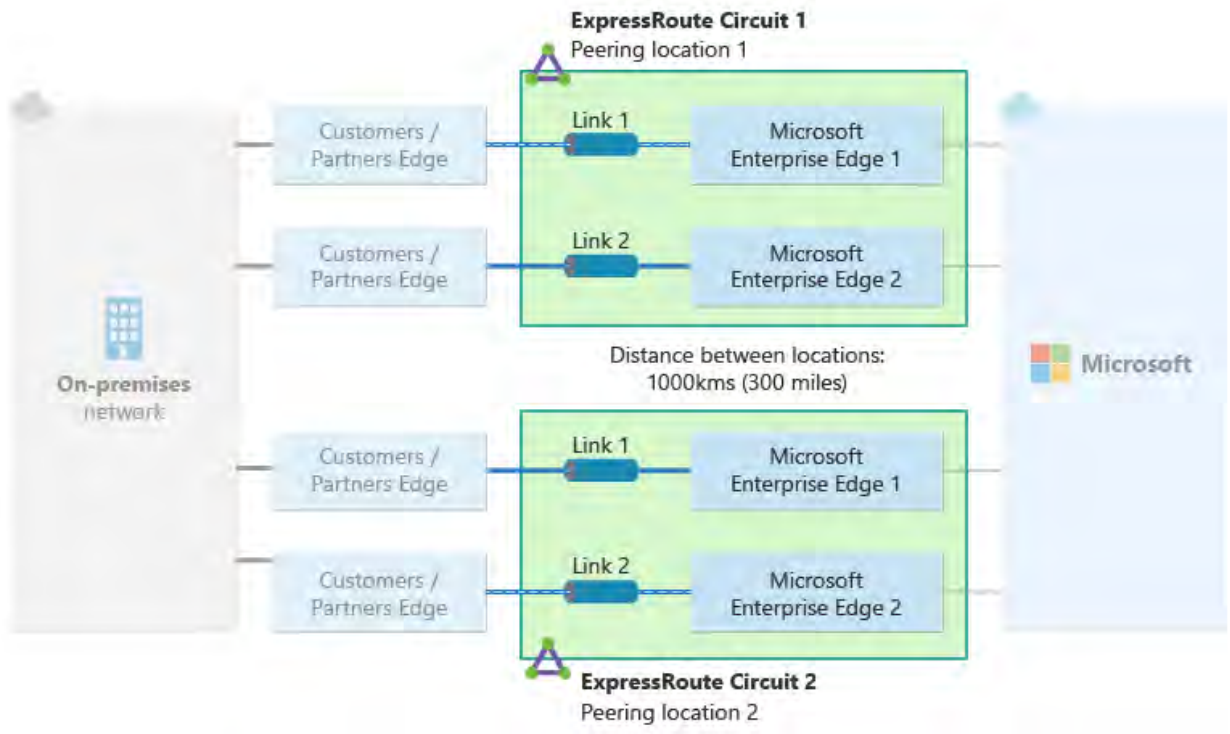

Achieving this goal requires observing two key aspects during the access network design:

- **Site Resiliency**: Ensuring no single point of failure exists at the network edge. Azure recommends a Metro configuration of 2+ sites configuration to maximize SLA based on site resiliency.
- **Zonal Resiliency**: Leveraging Azure regions and availability zones to maintain connectivity during localized failures.
  - **Availability Zones**: Deploy ExpressRoute Virtual Network Gateways as zone redundant. Availability zones provide fault isolation by spanning multiple physical locations within a region.
  - **Region-Level Resiliency**: Consider geo-redundancy by provisioning ExpressRoute circuits in multiple regions to guard against regional outages.

By adopting a multi-tiered approach that combines maximum or high resiliency architectures with zonal resiliency and robust monitoring, mission-critical applications can be safe-guarded against failures and outages.

More details on the SLA targets of such design can be found at the Azure specific SLA page: https://www.microsoft.com/licensing/docs/view/Service-Level-Agreements-SLA-for-Online-Services?lang=1&year=2025.

From an Aviatrix perspective, the multi-site best practice can be met by the following design decisions on how to contract and deploy Azure ExpressRoute:

- At least a pair of Active-Active Direct or hosted ExpressRoute circuits per location.
- Although the usage of another ERGW does not alter the ExpressRoute SLA per se, it would add additional protection against zone-level failures. Also, given the 1:1 ratio between VNet and ERGW, such approach would also require another Transit VNet to be deployed.
- ERGW SKUs should be carefully selected to match the bandwidth of the circuits in question. Planning for the usage of FastPath early on should also be considered when applicable.
- The hybrid DNS design needs to take into consideration the usage of FastPath, as that would enforce the usage of private resolvers in the hub VNet only (the Transit VNet hosting the ERGW).
- In Azure the allocation of ERGW in a VNet will require a /27 for the *GatewaySubnet* which should easily fit in the Aviatrix default VNet templates for FireNet (if a /23 is being considered).
- The Azure feature Global Reach could still be explored in case branch-to-branch connectivity is desired via ExpressRoute circuits and could also be leveraged by Edge gateways in Transit mode using Aviatrix Transit peerings. For more information on ExpressRoute Global Reach, please check the following page: https://learn.microsoft.com/en-us/azure/expressroute/expressroute-global-reach.

BGP Settings, BFD, ERGW and FastPath

While leveraging Aviatrix Edge overlay connectivity as part of the hybrid cloud connectivity, the underlay connectivity over Azure ExpressRoute needs to provide the basic connectivity between the main interface of the Aviatrix Transit gateways in the target VNet, whose attached ERGW is associated with the ER connection in question, and the on prem router granting the Aviatrix Edge gateway access to such environment.

This basically means two things:

- The ERGW will have to advertise only the CIDR prefixes that include the Transit gateways which the Edge gateways need to attach to. The ExpressRoute Gateway quota currently supports up to 1,000 prefixes, which is way more than enough assuming it can extend connectivity to up to 10 VNets by default or up to 100 with Premium add-on enabled.
- The BGP peer on the customer side will have to advertise at minimum the WAN subnets hosting all the Aviatrix Edge gateways that are supposed to attach to the aforementioned Aviatrix Transit gateways.

Unlike AWS, Azure does not have pre-defined BGP communities to influence how traffic gets load balanced towards the ExpressRoute circuits from the VNet side. Although, a similar approach is recommended with custom BGP communities being configured on the on-premises underlay BGP peer handling Edge WAN subnets with proper BGP local preference set, as the Azure private peering path selection will follow long-prefix match (LPM) logic. Such approach should promote better allocation of the overlay sessions across the Active-Active ExpressRoute connections.

For more information on how to use BGP setting to influence path selection on Azure ExpressRoute private peering, please check the following page: https://learn.microsoft.com/en-us/azure/expressroute/expressroute-optimize-routing.

Azure ExpressRoute supports BFD over private peering, reducing failure detection time over the Layer 2 network between Microsoft Enterprise Edge (MSEEs) and their BGP neighbors on the on-premises side from about 3 minutes (default) to less than a second.

For more information on how to configure BFD on most router vendors, please check the following Azure page: https://learn.microsoft.com/en-us/azure/expressroute/expressroute-bfd.

Finally, while the ERGW facilitates the exchange of network routes and directs network traffic, the ExpressRoute FastPath feature enhances data path performance between your on-premises network and your virtual networks. When enabled, ExpressRoute FastPath routes network traffic directly to virtual machines, bypassing the ExpressRoute virtual network gateway.

For more information on FastPath and corresponding ERGW SKU selection, please check the following Azure page: https://learn.microsoft.com/en-us/azure/expressroute/about-fastpath.

GCP Cloud Interconnect Best Practices for Max SLA Targets

Google recommends similar approach to take advantage of the better guaranteed SLA to maximize resiliency of a Cloud Interconnect deployment that will support mission-critical applications that have a low tolerance for downtime.

Multi-Site Cloud Interconnect Design

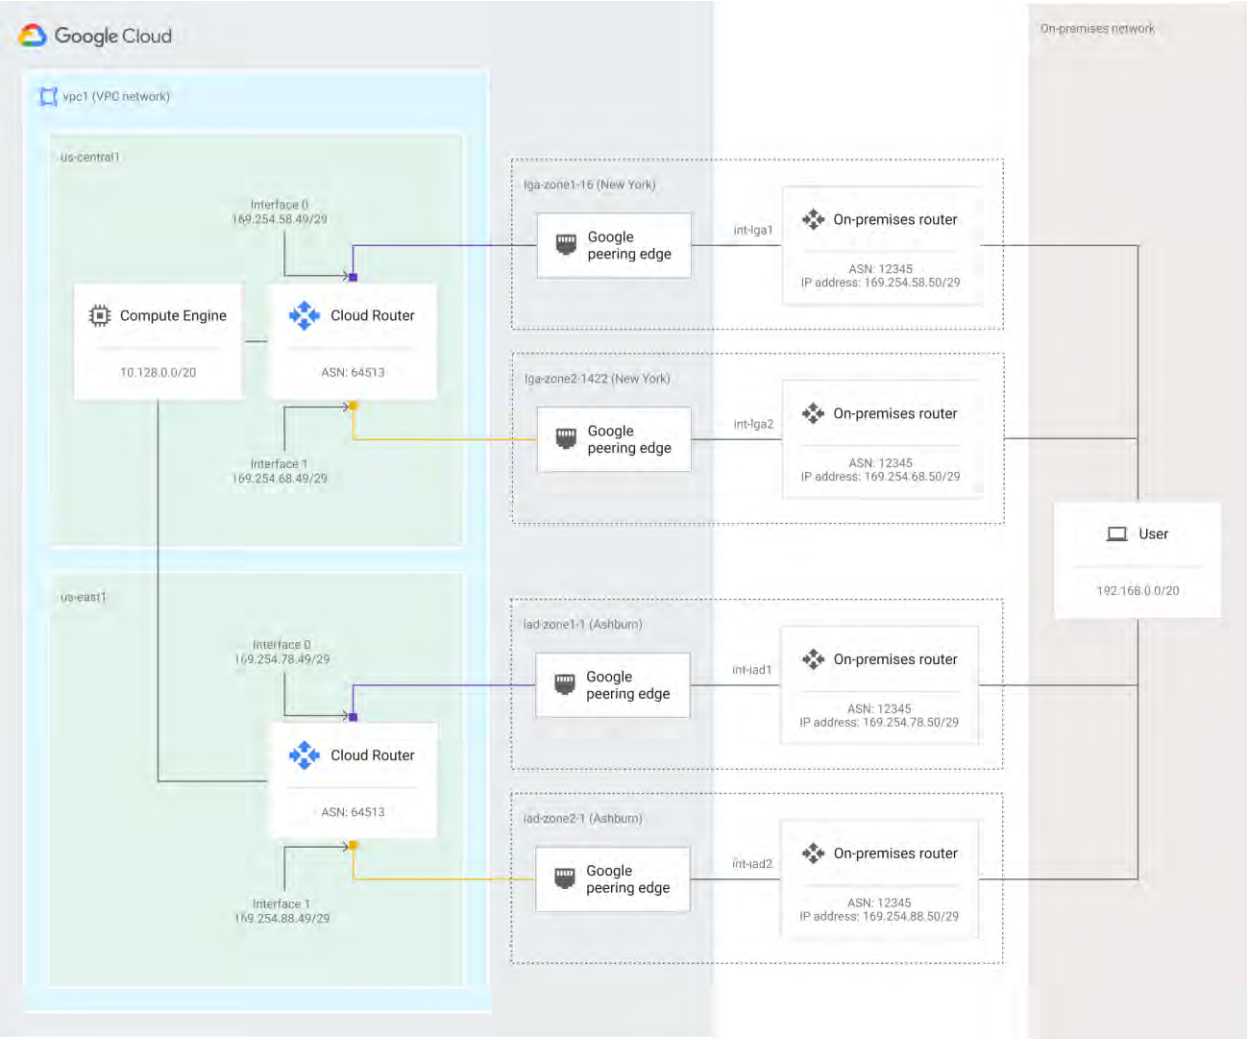

Google recommends the following best practices to achieve highest SLAs:

- At least 4x VLAN attachments, 2x per Google Cloud region. Each pair of VLAN attachments must have its own Cloud Router (two different Cloud Routers).

- Even though only two Cloud Routers are required, topologies with four Cloud Routers, one for each VLAN attachment (each in different edge availability domains), also meet the SLA requirement. For more information on availability domains and other key concept, please check the following GCP page: https://cloud.google.com/network-connectivity/docs/interconnect/concepts/terminology#locations
- The VLAN attachments in one region must connect to a Partner Interconnect connection in one metropolitan area (metro), and attachments in the other region must connect to a connection in another metro.
- Each VLAN attachment must be assigned to a different connection and to a different edge availability domain. This is because maintenance windows aren't coordinated across different metros, but they are coordinated for different edge availability domains within the same metro.
- Google recommends enabling VPC dynamic routing mode to global, so Cloud Router can advertise all subnets and propagate learned routes to all subnets regardless of the subnet's region. Although, this is not relevant when leveraging Aviatrix regional Transit VPCs. It would also be undesirable to have sub-optimal paths to reach out to Aviatrix Transit gateways (since Cloud Routers would announce remote VPC subnets with lower priority/MED).

More details on the SLA targets of such design can be found at the GCP specific SLA page: https://cloud.google.com/network-connectivity/docs/interconnect/sla.

From an Aviatrix perspective, the multi-site best practice can be met by the following design decisions on how to contract and deploy GCP Cloud Interconnect:

- At least a pair of Active-Active dedicated or hosted/Partner Cloud Interconnect connections per site location.
- A pair of VLAN attachments should be associated with a Cloud Router in each Transit VPC, so the underlay can be extended to up to 8x VPCs. This would correspond to the number of Transit VPCs maintained by Aviatrix globally considering the same connection. For more information on such quotas, please check the GCP specific quota page: https://cloud.google.com/network-connectivity/docs/interconnect/quotas.
- This setup allows for each site location to connect directly to any regions in GCP, as long such VPCs have their Cloud Routers associated with a pair of VLAN attachments (each on different edge availability domain).
- If more than 8x Transit VPCs are required by design, whether to accommodate for more VPCs within a hub in a region, or because more than 8x regions are being

considered, then the GCP quota of "*Peerings per VPC network*" needs to be considered as its default value is only 25. For more information on such GCP quota, please check the following GCP page: https://cloud.google.com/vpc/docs/quota#per_network.

- If adjusting the "*Peerings per VPC network*" quota is not a possibility, then other design mechanisms can be considered:
    - Leverage Aviatrix Multi-Tier Transit feature to extend the overlay connectivity to another Transit VPC within the same region.
    - Leverage VPC Network peering to extend the Cloud Interconnect VLAN attachments connectivity to up to another 25x Transit VPCs not necessarily in the same region, but that needs access to on premises via the same connections/VLAN attachments. This setup should still guarantee bandwidth benchmarks as Network Peering destination are considered as "within" VPC boundaries by GCP.
- The GCP Cross-Site Interconnect could still be explored in case branch-to-branch connectivity is desired via Cloud Interconnect connections and could also be leveraged by Edge gateways in Transit mode using Aviatrix Transit peerings. For more information on GCP Cross-Site Interconnect, please check the following page: https://cloud.google.com/network-connectivity/docs/interconnect/concepts/cross-site-overview.

BGP Settings, BFD and Cloud Routers

GCP Cloud Router is more flexible than its other CSPs equivalents as it tries to follow the RFC-4271. It supports the usual features of path selection based on LPM, BGP communities and it goes beyond with more advanced features, such as inbound and outbound filtering.

The Cloud Router offers two modes for best path selection: standard and legacy. The best path selection mode applies to all learned routes through Cloud Router in all regions of a VPC network, including custom learned routes. While both path selection modes would work while maintaining the Cloud Interconnect route exchange as Aviatrix underlay, it is recommended to leverage the standard mode, as influencing preferred paths from on premises is desirable, so a proper distribution of overlay session can be guaranteed across all the VLAN paths considered in case the admin does not want to rely solely on automatic ECMP.

This basically means two things:

- Each Cloud Router will have to advertise only the CIDR prefixes that include the Transit gateways which the Edge gateways need to attach to. The Cloud Router quota supports up to 200 custom prefixes, but that should not be needed as default advertisement of local subnet ranges suffices.
- The BGP peer on the customer side will have to advertise at minimum the WAN subnets hosting all the Aviatrix Edge gateways that are supposed to attach to the aforementioned Aviatrix Transit gateways.
- As part of the On-Ramp planning, if more than 8x Transit VPCs are required to follow the high availability best practices for GCP, then we also must consider:
    - Whether to have Cloud Routers distributed in pairs in separated VPCs from the Aviatrix Transit gateways that are then connected via VPC peering instead.
    - Whether to elect an (existing) Aviatrix Transit gateway as Multi-Tier, so the hub capacity for the selected regions can be extended to other Aviatrix Transits via Aviatrix Transit peerings.

Unlike AWS, GCP also does not have pre-defined BGP communities to influence how traffic gets load balanced towards the Cloud Interconnect circuits from the VPC side. Although, a similar approach is recommended with custom BGP communities being configured on the on-premises underlay BGP peer handling Edge WAN subnets with proper BGP local preference set, as the GCP Cloud Router standard path selection will follow long-prefix match (LPM) logic. Such approach should promote better allocation of the overlay sessions across the Active-Active Cloud Interconnect connections if the admin does not want to rely solely on the GCP automatic ECMP.

For more information on how to use BGP settings and routing policies to influence path selection on GCP Cloud Routers serving Interconnect attachments, please check the following pages:

- https://cloud.google.com/network-connectivity/docs/router/concepts/learned-routes#dynamic-routing-mode-effects-on-learned-routes
- https://cloud.google.com/network-connectivity/docs/router/concepts/bgp-route-policies-overview.

While connecting GCP with on-premises networks with Dedicated Interconnect or Partner Interconnect, enabling BFD for fast detection of link failure and failover of traffic to an alternate link that has a backup BGP session is recommended to maintain high-availability network connectivity paths that can respond quickly to link failures.

For more information on how to configure asynchronous BFD in GCP Cloud Router, please check the following GCP documentation: https://cloud.google.com/network-connectivity/docs/router/concepts/bfd.

Aviatrix Edge Best Practices

Now that each CSP best practices have been described and correlated to an Aviatrix design pattern to facilitate both site as well as link redundancy to maximize SLAs, the next step is to observe the specifics of the Aviatrix Edge gateway deployment itself, observing the best way to establish L1/L2/L3 connectivity on both WAN, LAN and management interfaces, while determining best sizing, routing mode, specific BGP settings, route filters and policies.

The best practices documented in the following section should be valid to any Aviatrix Edge gateway regardless of the selected platform (self-managed, AEP or Equinix/Megaport), otherwise the exceptions will be highlighted when applicable.

Aviatrix Edge Interfaces

L1/L2 Connectivity

By default, the Edge gateways will have 3x interfaces, either physical (AEP) or virtual (self-managed, Equinix/Megaport).

From a L1 aspect, which is applicable only to AEP platform, the selected media to connect each interface will depend on the network cards selected and currently approved hardware.

For more information on currently supported hardware, please check the following Docs page: https://docs.aviatrix.com/documentation/latest/network/edge-hardware-specs.html.

Once the NIC is identified, please use the NIC specification to check for compatible peripherals and optics to use with each NIC in the vendor website.

For example:

- AEP running on Dell R450 using NICs Intel E810-XXV.
- Check on vendor page https://compatibleproducts.intel.com.
- Confirm that SFP28 SR Optics with Dell SKU 407-BCBG is supported.

From a L2 aspect, the usage of VLAN is recommended when deploying Edge in AEP just as a general good networking best practice. The same is valid for ESXi and KVM environments, although VLAN 802.1q tagging is often not required and overlooked.

At Equinix Network Edge, the Aviatrix Edge devices will leverage VRF/L3 separation between its interfaces and sub-interfaces/VLAN tagging is not required. For more information, please check the Aviatrix specific page in Equinix Docs: https://docs.equinix.com/network-edge/vendors-devices/aviatrix/ne-aviatrix-specs.

At Megaport Virtual Edge, Megaport will automatically tag the Aviatrix Edge interfaces unless configured to do otherwise. VLAN tagging in Megaport follows 802.1ad by default which is not supported by Aviatrix, so Megaport will not include the outer tag on traffic handled by the WAN interface towards the VXC or fabric. For more information on Aviatrix as MVE, please check the following Megaport Docs page: https://docs.megaport.com/mve/aviatrix/.

Management Connectivity

The network interface to connect to the Aviatrix Controller is the management interface, often referred as eth2, but aliases might vary depending on specific port mapping. This interface requires a default gateway, DNS access and direct access to the Aviatrix Controller, for Aviatrix software downloads, and tracelog uploads.

Internet access is not required, if Management interface is over a private network. The management traffic can also be proxied as needed.

Statically assigning the IP configuration to the management interface can be done on both self-managed as well as the AEP platform supported hardware, while DHCP is preferred for both Equinix Network Edge and Megaport MVE. Equinix Network Edge will also accept static IP assignment if "*over private network*" mode is selected.

Please note that the AEP platform will use the default CIDR 10.100.101.0/24 as an internal bridge for the outer management interface. This subnet address can be changed during initial onboarding steps as needed.

WAN Connectivity

The network interface to connect to the Aviatrix Transit Gateway is the WAN interface, often referred as eth0, but aliases might vary depending on specific port mapping. This interface requires a default gateway and L3 reachability to Aviatrix Transit Gateways Private or Public IPs.

The WAN interface can follow distinct profiles depending on the platform considered:

- The self-managed Edge gateways currently support single WAN interfaces only (Spoke mode only as of Controller version 8.0).
- The AEP platform will depend on the NIC configuration, and it currently supports up to 4x WAN interfaces in Transit mode. Current profiles can be selected as part CoPilot onboarding process during device configuration.
  For more up to date information, please check the following page:
  https://docs.aviatrix.com/documentation/latest/network/edge-device-os-idrac.html#2-onboard-edge-device-onto-aviatrix-edge-platform.
- Edge gateways at Equinix Network Edge support up to 9x WAN interfaces per gateway when in Transit mode.
- Edge gateways as Megaport MVEs support up to 4x WAN interfaces per gateway when in Transit mode.

LAN Connectivity

The network interface to connect to BGP LAN peers is the LAN interface, often referred as eth1, but aliases might vary depending on specific port mapping. This interface requires L2 connectivity to the BGP LAN peers and it is often configured as part of /31 networks.

The LAN interface is used only when the Aviatrix Edge is configured in Spoke mode. DHCP mode is not support in Equinix, Megaport and AEP.

Static IP assignment on LAN interfaces is supported in all platforms. Extra steps might be required if selecting Active-Standby mode as VRRP settings are also required.

Site Redundancy Best Practices

Site redundancy best practices will primarily follow the recommended best practices established by each one of the CSPs, which in practical terms should guide the Aviatrix Edge design towards the following:

- No single CSP private link failure should cause the disconnection of any Edge gateways to each respective remote WAN attachments. So, to address this risk, the WAN underlay connectivity for the Edge should allow the L3 connectivity between Edge gateways and its destination over all the available paths. Using BGP attributes to control load balancing is possible and recommended. BFD is a hard requirement

to prevent a single private link/connection failure to also trigger an unnecessary overlay failover.

- No single switch failure in the site in which the Edge gateway is racked or virtualized should disconnect it completely from the on-premises destinations. This means that regular DC network best practices need to be followed and Edge gateways should establish cross-BGP sessions with the stack/redundant L3 switches providing LAN connectivity while VIPs (VRRP/HSRP) can be maintained as next-hops.
- No single Edge gateway failure should disconnect an entire site from its cloud destinations. This means that Edge gateways should be distributed equally across two or more geo-locations when applicable (COLO sites or DCs) and they should always be deployed with at least N+1 Active-Active redundancy.

A good example of how this setup should look like from a high-level design perspective when GCP Cloud Interconnect is considered with two COLO locations being leveraged for On-Ramp access:

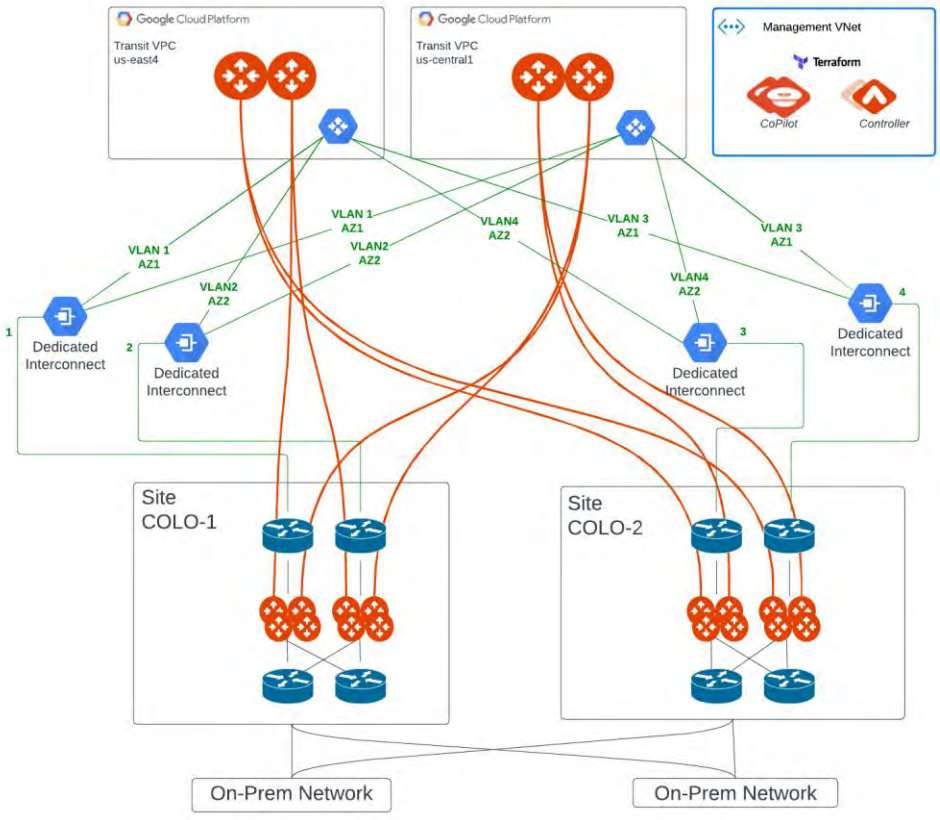

Capacity Planning and Sizing

The capacity planning when deploying Aviatrix Edge gateways start with the size of each Edge unit itself. When deploying the Edge platform that leverages compute version of Edge, the following t-shirt sizes based on number of CPU cores can be leveraged. Please note that the memory profile might vary per platform:

Size	Hardware Profile	Storage Requirements	Single Node Benchmark
Small	2 vCPU - 4GB	64 GB	< 1Gbps throughput
Medium	4 vCPU - 8GB	64 GB	< 5Gbps throughput
Large	8 vCPU - 16GB	64 GB	~10Gbps throughput
X-Large	16 vCPU - 32GB	64 GB	> 10Gbps throughput

These profiles are valid for self-managed VMs (ESXi and KVM), Equinix Network Edge and Megaport as well. Megaport *x-large* size only supports 12x vCPUs.

The Aviatrix Edge Platform sizing is directly tied to approved hardware specifications, which can be checked in the following Docs page: https://docs.aviatrix.com/documentation/latest/network/edge-hardware-specs.html.

Currently, each R450 server can benchmark up to 25Gbps when jumbo frames are being considered. All other listed bandwidth benchmarks might require optimal conditions, such as the usage of jumbo frames.

If bandwidth requirements exceed the bandwidth benchmark of a single node at maximum size in the selected platform, then horizontal scaling should be considered as well as Active-Active mode, for example:

- Current requirement for bandwidth is aggregated 100Gbps with 2x A/A 50Gbps links in the referred On-Ramp access network.
- Selected platform is AEP and top listed hardware model is Dell 450 using Intel E810 cards, each node capable of 25Gbps.
- At least, 4x Edge gateways should be considered and with a 2x2 distribution in case the private links are also terminated on different co-locations.
- Downstream L3 switches peering with such Edge gateways should be configured to perform ECMP per flow to leverage the full aggregated bandwidth capacity as well.

Spoke or Transit Mode Selection

The Aviatrix Edge will support two types of mode just like regular gateways:

- **Spoke mode:** Requires the usage of LAN interface, does not support Transit peerings or Spoke to Spoke peerings. Only external connectivity allowed is BGPoLAN via its LAN interface. When in Active-Standby mode, it can also support static routes when Edges are configured as part of local VLANs.
- **Transit mode:** Not supported in self-managed platform as of Controller version 8.0. Supports Transit peerings and regular spoke attachments (of other Aviatrix gateways in cloud). Additional support to external connectivity includes BGPoIPsec and BGPoGRE, over WAN interfaces. Transit mode should be considered when:
    - Branch-to-branch connectivity is being considered via private connectivity using features like AWS SiteLink or Azure Global Reach.
    - When part of a solution pattern that includes DCF (Aviatrix Cloud Firewall) and Aviatrix FireNet is not being considered, so spokes can attach directly to Edge gateways.
    - Please note that Aviatrix FireNet can also be deployed with Aviatrix Edge Transit mode (EaT) as Transit peering is also supported.

Advanced Settings

As part of advanced tuning of the Aviatrix environment, the following features and attributes should be configured to favor performance and resiliency, as indicated in the list below:

- **Jumbo Frame:** It improves the performance throughput between an Aviatrix Transit Gateway and Edge Gateway. It is supported on all Edge platforms. There are two Jumbo Frame configuration settings for an Edge Gateway: one for the Edge Gateway and another when you create an Edge Gateway attachment. Jumbo Frame is fully supported in AWS and OCI. GCP has limited support (requires unencrypted VLAN attachment). Azure does not support it.
- **BGP timers:** Aviatrix allows access to the following global BGP timers, that will affect all BGP sessions within any gateway being considered:
    - **BGP Hold Time:** The supported range is 12 to 180 seconds.
    - **BGP Polling Time:** Aviatrix Transit and BGP-enabled Spoke Gateways report its BGP routes to the Aviatrix Controller periodically. This polling time affects

the BGP route change convergence time. The supported range is 10 to 50 seconds.
    - **BGP Neighbor Status Polling Time:** The supported range is 10 to 50 seconds.
- **Other attributes configured per connection:**
    - `prepend_as_path` : List of ASN to prepend AS path field. This argument supports both directions when considering Edge gateways in Spoke mode.
    - `spoke_bgp_manual_advertise_cidrs` : Manually advertise CIDRs over a specific BGP connection. This mode does not append CIDRs (aggregation) but replaces them with the configured list instead.
    - `enable_preserve_as_path` : This argument when enabled, preserves the AS path of manual advertisements to match the current active path for the CIDR which is maintained in the gateway local best path DB.
    - `bgp_bfd` : BGP BFD configuration applied to the specific BGP session. This setting does not affect the SDN routing over the Edge spoke attachment or Transit peering. Those scenarios are still covered by gateway and tunnel failures timers instead. BFD can be further configured as the following:
        - `transmit_interval` : BFD transmit interval in ms. Valid values between 10 to 60000. Default: 300.
        - `receive_interval` : BFD receive interval in ms. Valid values between 10 to 60000. Default: 300.
        - `multiplier` : BFD detection multiplier. Valid values between 2 to 255. Default: 3.
        - `enable_bgp_multihop` : Whether to enable multi-hop on a BFD connection.

- **RX Queue Size:** Performance tuning attribute. Allows for the following values: “1K”, “2K” and “4K”. Please contact Aviatrix support for proper guidance before adjusting this value in a production environment.
- **Keepalive Speed:** In normal state, Aviatrix gateways send keep alive messages to the Controller. Keep Alive Speed determines when Controller starts to evaluate if a gateway is down. Fast is recommended for most deployments and might require review of the Controller sizing depending on the number of gateways being managed.
- **Tunnel Down Detection Time:** Global setting accessible via CoPilot UI only. It can be configured via Terraform on a per-gateway basis. The supported range is 20 to 600 seconds.

Route Approval

Route Approval is another important advanced feature that allows the admin to filter which learned CIDRs on each Edge external connection should be approved to be added and distributed across the Aviatrix data plane.

There are two modes of Route approval:

- **Gateway mode:** As the name indicates it is a global setting that is applicable to all connections within the same gateway.
- **Connection mode:** It approves routes on a per connection basis. This is the preferred mode as it prevents the gateway from ever selecting a non-desirable path.

It is also important to note that, regardless of the mode approved, the Aviatrix gateway will only maintain one single active path for any CIDR at any given time, the exception to this rule is for paths with equal metric/cost, in which case the Aviatrix Controller will maintain and install all the paths, so the gateways can perform ECMP on the outgoing traffic flows.

Route approval can be handy when on premises advertises all CIDRs and uses route aggregation at the same time. This way the Aviatrix admin can choose to simply approve the summary ranges and possibly exclude default routes when desirable, to simplify and make the routing footprint more efficient.

For more information on how to deploy Route Approval on Aviatrix gateway in general, please check the following Docs page: https://docs.aviatrix.com/documentation/latest/network/bgp-route-approval.html.

Network Segmentation

Aviatrix Multicloud Transit Network Segmentation provides network isolation and better security. It does this through network/routing domains and connection policies. The segmentation is enforced globally at every Aviatrix Transit gateway. You can have up to 200 unique network domains on each Aviatrix Transit Gateway.

Besides the default approach of enforce L3 separation amongst the multiple app domains that might coexist within the same Transit hub/fabric in Aviatrix data plane, the network segmentation feature can become handy in Aviatrix Edge deployments when multiple sites

are connected, via different ASNs and Site IDs (on the Edge gateways themselves) while attached/connected to the same Aviatrix Transit gateways.

By default, ActiveMesh will distribute routes to Edge attachments when they do not share the ASN (and Site ID), for a scenario in which different COLOs provide connectivity to the same data center infrastructure. This is most of the time an undesirable situation and can easily be addressed by the adoption of Aviatrix network domains, one for each Site ID while not configuring a connection policy to allow such connectivity. This approach would prevent and cease any attempt of distributing CIDRs from one Edge location to another.

For more information on how to leverage and deploy network segmentation with Aviatrix, please follow up on the following Docs page: https://docs.aviatrix.com/documentation/latest/security/network-segmentation-secured.html.

For more information on how Aviatrix ActiveMesh routing technology works, please check the following Docs page: https://docs.aviatrix.com/documentation/latest/network/activemesh-about.html.

DCF (Distributed Cloud Firewall) Support

Distributed Cloud Firewall (DCF) uses micro-segmentation to provide granular network security rules for distributed applications in cloud environments. It enables network policy enforcement between SmartGroups, WebGroups, and ExternalGroups defined in a single cloud or across multiple clouds.

When compared to Network Segmentation, DCF can be seen as proper dynamic firewall access lists enforced at Spoke gateway level, while Network Segmentation promotes VRF-like separation at Transit gateways.

This distinction is important as it reinforces a possible hybrid strategy when a full at-scale DCF enforcement is not required or still beyond the current maturity of the cloud environment design.

DCF while enforced at the edge of the network (Spoke and Edge gateways) could be used together with Aviatrix FireNet and prevent unauthorized traffic to traverse to the Transit layer, to avoid consuming extra CPU cycles from both Transit gateway and firewall instances.

Another use case is to enforce micro-segmentation on traffic sourced from on premises at the first Aviatrix node (Edge gateways), for scenarios in which the traffic would be allowed by

route approval and network segmentation. In such scenarios, this traffic pattern would only be blocked by proper firewall rules in FireNet firewalls, after having reached to the Aviatrix Transit gateways.

Up to Controller version 8.0, Aviatrix Edge supports the basic version of DCF while leveraging mostly L4-based Smart Groups (IPs/CIDRs). The usage of this feature is not recommended in environments also leveraging Aviatrix gateways to perform NAT to solve for overlapping IP scenarios.

For more information on usage of DCF on regular spoke gateways, please check the following Docs page: https://docs.aviatrix.com/documentation/latest/security/dcf-overview.html.

NAT Support

Aviatrix Edge supports customized SNAT and DNAT for use cases where the CSP network CIDR overlaps with the on-prem network CIDR.

The following NAT scenarios are supported:

- **Single IP SNAT:** For network traffic initiated from the Edge location towards the CSP.
- **Customized SNAT:** For network traffic initiated from the Edge location towards the CSP. This method is not supported when VLAN segmentation is also configured in the same network domain (EaS in Active-Standby with VRRP).
- **Customized DNAT:** For network traffic initiated from Edge location towards Transit Gateway or CSP.

NAT over HPE connections is also not supported in Edge gateways as of Controller version 8.0. The usage of BGPoLAN interfaces is also not supported in attached Transit gateways, in both SNAT and DNAT configurations.

As a best practice, it is still recommended to perform both customized SNAT and DNAT on the Spoke gateways fronting the CSP workloads, for a better distribution of the NAT rules for environments in which overlapping needs to be resolved globally and at-scale. NAT over HPE is also not yet supported in regular spoke gateways as of Controller version 8.0.

For more information on how to perform Aviatrix Customized SNAT/DNAT on Edge gateways, please check the following Docs page: https://docs.aviatrix.com/documentation/latest/network/edge-spoke-snat-dnat.html.

For more information on how to perform SNAT/DNAT on regular spoke gateways, please check the following Docs page: https://docs.aviatrix.com/documentation/latest/network/spoke-gateway-snat-dnat.html.

Deployment Strategies

The installation procedures for Aviatrix Edge gateways will vary considerably, based on the selected platform. Please check the following section for documentation references for each platform installation guide.

Plan Aviatrix Edge Environment

1. Select the CSP reference architecture for highly available access network.
2. Select Edge platform. Start procurement of hardware as soon as possible, if choosing AEP.
3. Select Aviatrix redundancy mode to match CSP architecture of choice and current on premises/co-location switches setup: Active-Active or Active-Standby.
4. Select the Edge mode: Spoke or Transit.
5. Estimate initial capacity for Edge layer based on item #1 and application SLOs. This early decision is especially important when selecting AEP as it impacts the initial procurement for required hardware.
6. Plan for long term: how to horizontally-scale Aviatrix Edge gateways if needed when needed. Contact Aviatrix Product team for more details on our roadmap.
7. How to monitor Edge: besides CoPilot, will any other 3rd party tool consume Syslog or CoPilot API to continuously monitor Edge gateways capacity and links utilization as part of day-2 operational processes?

Basic Configuration Steps

1. Self-Managed in ESXi/KVM environments: https://docs.aviatrix.com/documentation/latest/network/edge-selfmanaged.html.
2. Aviatrix Edge as Equinix Network Edge device: https://docs.aviatrix.com/documentation/latest/network/edge-equinix.html.
3. Aviatrix Edge as Megaport Virtual Edge device: https://docs.aviatrix.com/documentation/latest/network/edge-megaport.html.
4. Aviatrix Edge Platform on supported hardware: https://docs.aviatrix.com/documentation/latest/network/edge-aep.html.

Advanced Configuration Steps

1. Build high-level design based on design decisions.
2. Document which Aviatrix advanced settings will be required if any. For example: whether Network segmentation is required or not, whether any AS prepending is required or not to facilitate any defined routing policy.
3. Plan for the underlay BGP configuration to favor failover efficiency with BFD and prevent unnecessary overlay failovers.
4. Double-check that underlay BGP settings are indeed following the underlay redundancy of choice. If Active-Active, the overlay sessions need to be properly balanced across the referred redundant links.
5. Review finalized design to confirm both Aviatrix and the CSP best practices have been followed.
6. Perform the validation steps documented in the next section to document failover duration baselines and benchmarks to be used for future reference. The validation plan should include both data plane as well as control plane events. The latter will often include disaster recover procedures to restore backups to recover from failed Aviatrix Controller instances.
7. Verify selected monitoring tools/processes recorded and alerted on all validated events as part of the failover testing.

For assisted help while installing and deploying Edge gateways, please reach out to your Aviatrix representative about our professional services offerings: https://aviatrix.com/platform/advanced-cloud-services/.

Design Validation

After implementing any solution, thorough testing is needed to ensure correct design implementation and requirement fulfillment.

Validation of the design entails testing the scenarios that are specific to the Aviatrix Edge On-Ramp environment. Such tests must include failover scenarios affecting all connectivity that facilitates the On-Ramp design, so environment and CSP specific benchmarking is established early on and documented as evidence in the low-level document being considered.

Example of Test Cases for Validation Runbook

#	Use Case	What to do	What to see	Expected behavior	Related features
1	Prig from cloud VM to on-prem destination	Ping/traceroute/curl private IP of on-prem destination	Connection established and L3 path via connected Transit (Edge gateways included)	Both ping and traceroute are successful. Curl test should confirm no asymmetry issues in the path.	Validates any firewall rules and network segmentation policies in the path
2	Failover of Transit gateway	Ping/traceroute/curl destination private IP and then stop the Transit gateway being used by current flows	Time to recovery of the current flow when applicable	Both ping and traceroute recover. Curl test should still confirm no asymmetry issues in the path.	Controller monitoring of gateway status and response to network events to update routes (as needed).
3	Failover of Edge gateway	Ping/traceroute/curl destination private IP and then stop the Edge gateway being used by current flows	Time to recovery of the current flow when applicable	Both ping and traceroute recover. Curl test should still confirm no asymmetry issues in the path.	Controller monitoring of gateway status and response to network events to update BGP routes (as needed).
4	Failover of entire Site ID (only applicable if site failover is required)	Ping/traceroute/curl destination private IP and then stop ALL Edge gateways sharing the same Site ID being used as next hop by current flows	Time to recovery of the current flow when applicable	Both ping and traceroute recover. Curl test should still confirm no asymmetry issues in the path.	Controller monitoring of gateway status and response to network events to update BGP routes, as well as underlay connectivity between the considered on-prem sites.
5	Underlay failover (single private link, either DX, ER or IC)	Ping/traceroute/curl private IP of on-prem destination after shutting down one of the private links connections	Connection established and L3 path via connected Transit (Edge gateways included). Edge gateways overlay will reestablish via remaining local private links. Failover should be seamless if	Both ping and traceroute are successful. Curl test should confirm no asymmetry issues in the path.	Controller monitoring of overlay connectivity should not detect the underlay failover to local private link underlay connectivity is properly configured.

			underlay connectivity is properly configured.		
6	Underlay failover (all private links towards a CSP within the same site/co-location)	Ping/traceroute/curl private IP of on-prem destination after shutting down all the private links that connects a site/co-location to a CSP	Connection established and L3 path via connected Transit (Edge gateways included from secondary site/co-location).	Both ping and traceroute are successful. Curl test should confirm no asymmetry issues in the path.	Controller monitoring of overlay connectivity should detect the underlay failure and withdrawn all routes from such CSP on the connected Edge gateways. Edge gateways will them withdrawn corresponding paths from LAN peers.
7	Edge LAN BGP single peer failure	Ping/traceroute/curl private IP of on-prem destination after shutting down all the BGP sessions to the same LAN peer	Connection established and L3 path via connected Transit. Requires LAN cross-connectivity with site/co-location switches to avoid failover event.	Both ping and traceroute are successful. Curl test should confirm no asymmetry issues in the path.	Edge local detection of link status should maintain the active/valid paths for the active LAN peers. BFD can improve this scenario if properly configured.
8	Edge LAN BGP all peers failure	Ping/traceroute/curl private IP of on-prem destination after shutting down all LAN sessions of one or more Edge gateways in the same Site ID	Connection established and L3 path via connected Transit via Edges still advertising the same destination.	Both ping and traceroute are successful. Curl test should confirm no asymmetry issues in the path.	Controller monitoring of overlay connectivity should detect all paths were withdrawn to affected Edge gateway(s) and update the route advertisement on the respective attachments.

9	Controller failover (restore from backup without IP migration)	Ping/traceroute/curl private IP of on-prem destination	Connection established and L3 path via connected Transit (Edge gateways included)	Both ping and traceroute are successful. Curl test should confirm no asymmetry issues in the path.	The Controller restore from backup should NOT impact the data plane.
10	Controller failover (restore from backup with IP migration)	Ping/traceroute/curl private IP of on-prem destination	Connection established and L3 path via connected Transit (Edge gateways included)	Both ping and traceroute are successful. Curl test should confirm no asymmetry issues in the path.	The Controller restore from backup should NOT impact the data plane. The admin will have 24h to locally update the Controller IP on each Edge gateway.

Monitoring and Management

The following Aviatrix monitoring tools are applicable and fully aligned with the Aviatrix Edge On-Ramp design:

- Aviatrix CoPilot Notifications with Webhooks integrated with paging or ticketing systems.
- Aviatrix CoPilot APIs (Status and Metric) using Prometheus framework integrated with monitoring systems like Grafana and other alerting systems.
- Exported Syslog streams from gateways and Controllers.

Once the day-2 operational monitoring processes have been defined, solid naming conventions of all resources will facilitate for more efficient processes. More information on best practices is included in the next section.

The Aviatrix Status and Metric APIs could feed a system for constant monitoring of not just the status and system utilization, but in parallel to other native monitoring to evaluate private link utilization and overall load balancing.

The overall Aviatrix Day-2 Operations best practices are documented in the following Docs page: https://docs.aviatrix.com/documentation/latest/monitoring/index.html.

Glossary

Please refer to the Aviatrix Docs page for a complete version of this glossary at:

- https://docs.aviatrix.com/documentation/latest/getting-started/platform-overview/general-glossary.html.
- https://docs.aviatrix.com/documentation/latest/getting-started/platform-overview/aviatrix-glossary.html.

Term	Definition
ACL	Access List
API	Application Programmable Interfaces
ASN	Autonomous System Number
BFD	Bidirectional Forwarding Detection
BGP	Border Gateway Protocol
CIDR	Classless Inter-Domain Routing
CSP	Cloud Service Provider
Cloud On-Ramp	Direct private connectivity to multiple CSPs
DHCP	Dynamic Host Configuration Protocol
DNAT	Destination Network Address Translation
DNS	Domain Name System
DX	AWS Direct Connect
ECMP	Equal-Cost Multi-Path routing
E-W	East West Traffic
EOL	End of Life
ER	Azure ExpressRoute
FW	Firewall
GW	Gateway
HA	High Availability
HPE	High Performance Encryption, Aviatrix IPsec Tunnel Bundle Technology
Hub	Transit location of shared services and network core
HTTPS	Hypertext Transfer Protocol Secure
HSRP	Hot Standby Router Protocol
IC	GCP Cloud Interconnect
MED	BGP Multi-Exit Discriminator
NIC	Network Interface Card
RX	Receive

Term	Definition
SFP	Small Form-factor Pluggable
SKU	Stock Keeping Unit
SLA	Service Level Agreement
SLO	Service Level Objective
SNAT	Source Network Address Translation
VLAN	Virtual Local Area Network
VRRP	Virtual Router Redundancy Protocol