

Arista CloudEOS Multi-Cloud Solution Guide

Enabling enterprise networking 'at the speed of Cloud'

Table of contents

Why Multi-Cloud?	3
Multi-Cloud Architecture and Challenges	3
Arista CloudEOS - Multi-Cloud Networking Solution	4
Solution Goals	4
Solution Components	5
1 - CloudEOS Router VM	6
<i>Underlay network in the public cloud</i>	7
<i>Overlay network in the public cloud</i>	7
2 - CloudEOS Router for Kubernetes	10
3 - Hashicorp Terraform Provider	10
4 - Cloudvision	11
Putting it all together	13
CloudEOS Use-Cases	14
Secure Multi-Cloud Edge	15
Multi-Cloud Path Optimization	18
Consistent Segmentation with Central Policy Enforcement	19
<i>Multi-Cloud Network Segmentation</i>	20
<i>Firewall Insertion for Central Policy Enforcement</i>	21
CloudEOS and AWS Transit Gateway Integration	22
Visibility and Governance	24
Summary	29
Appendix	29

Why Multi-Cloud?

With more employees now working from home and customers accessing information and services from their mobile devices, organizations are accelerating their cloud migration, and multi-cloud adoption for fast application delivery and high velocity innovation. Organizations like the flexibility provided by multiple public cloud service providers and are looking to optimize workloads as they migrate, in addition to managing cost and introducing better governance to ensure operational efficiency.

Multi-Cloud Architecture and Challenges

Multi-cloud architecture is the logical next step in the evolution that allows organizations to distribute their workloads across the various cloud services providers. Some benefits the organizations see with this approach are:

- Improved application availability with reduced latency
- Superior security while meeting regulatory requirements
- Optimized ROI (Return On Investment)
- Autonomy by avoiding vendor-lock in
- Higher reliability

The previous decade was about moving all application workloads to the public cloud, however in the past few years, most organizations have now adopted a hybrid multi-cloud strategy wherein they maintain both an on-premise private cloud and leverage the flexibility and agility provided by the public cloud. This is done primarily for security, compliance, data protection and autonomy / prevent vendor-lock in.

Multi-Cloud architecture has its own challenges as well. The common ones are:

- **Inconsistency and non-repeatable:** each cloud provider has its own uniqueness like disparate network architectures, features, and scales which create a steep learning curve for customers to operate in the cloud and create operational challenges across existing environments like data center and campus networks. Thus, bringing up a new VPC, VNET and connecting it to the enterprise's existing environment usually requires weeks of work.
- **Limited visibility and observability:** most enterprise customers find it extremely difficult to troubleshoot a network issue in the public cloud due to the lack of information and visibility, especially when troubleshooting requires cross domain coordination and packet-level observability.
- **High cost:** Oftentimes, enterprise customers find that public cloud doesn't always reduce IT cost. Different applications sharing data across different clouds will increase data charge dramatically. Adopting a new cloud provider usually means adding 4~6 engineers to support production services, different runbooks, ongoing network changes, escalations, which would impact the organization's overall budget plan.

Arista CloudEOS - Multi-Cloud Networking Solution

CloudEOS is Arista's multi-cloud and cloud-native networking solution supporting autonomic operation to deliver an enterprise-class, highly-secure, and reliable networking experience for extending an enterprise network to any cloud. As part of the Arista EOS and CloudVision product family, it delivers consistent segmentation, automation, telemetry, provisioning and troubleshooting for the enterprise edge, WAN, campus, data center and multiple public and private clouds.

To provide a scalable and automated network experience, CloudEOS integrates with Arista CloudVision to simplify the operators experience of interconnecting and managing multi-cloud, cloud-native and on-premises enterprise networks. Leveraging a network-wide approach for workload orchestration and workflow automation, together with advanced network telemetry and standards-based programmability, CloudEOS and CloudVision provide a seamless and universal approach to multi-cloud networking.









Figure 1: Arista CloudEOS Solution Overview

Solution Goals

In a multi-cloud deployment, flexibility and agility in managing the solution are key:

- Launching an app into any cloud while maintaining network consistency
- Easily scaling up across new regions and providers
- Replatforming to a new cloud based on geography, performance, or cost
- Re-entry back into an on-premises facility if that is the best business decision

To support these goals the following technology components are central to Arista’s quest in providing a uniform multi-cloud solution:

 <p>Declarative Provisioning and Deployment</p>	<p>Single-click deployment of multi-cloud networks in minutes. Single tool for management through real-time state streaming telemetry and network topology.</p>
 <p>Global Multi-Cloud Network</p>	<p>Build a global secure multi-cloud network through a consistent cloud architecture and operational model.</p>
 <p>Dynamic Path Selection</p>	<p>Business SLAs delivered over the best path in the cloud network, with rich visibility and fine-grained control.</p>
 <p>Consistent Network Segmentation</p>	<p>Standards-based network segmentation model that scales and extends the reach of network trust zones, using best-in-class firewalls to enforce policies for multi-cloud workloads.</p>
 <p>Network Elasticity</p>	<p>Scale the network up and down, in and out to meet business needs and lower OpEx.</p>

Solution Components

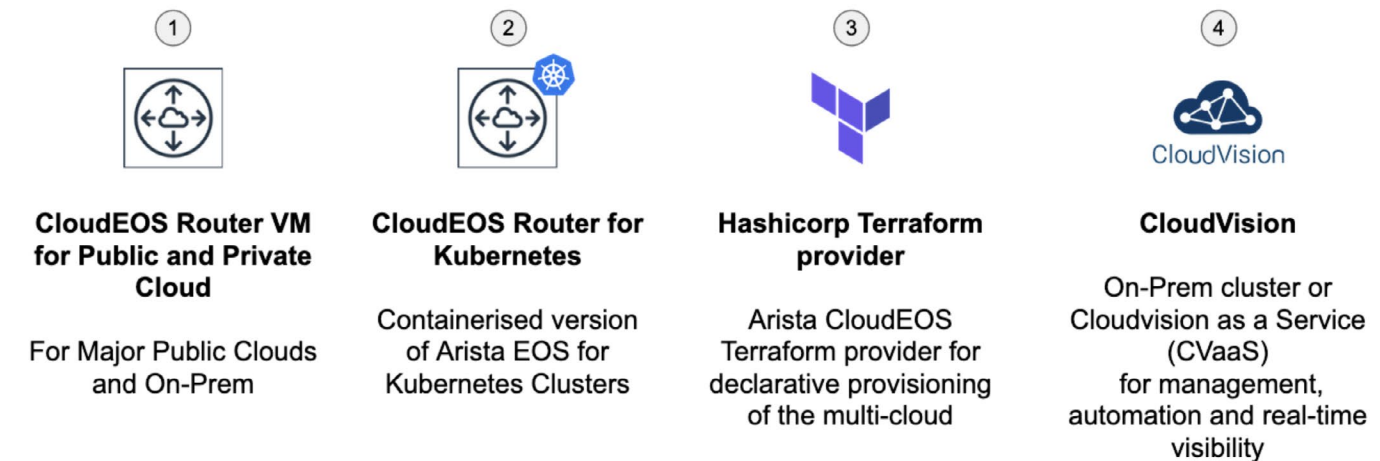


Figure 2: Arista CloudEOS Solution Components

1 - CloudEOS Router VM

The [Arista CloudEOS™ Router VM](#) is a cloud-grade, feature-rich, multi-cloud and multi-hypervisor virtual router that enables enterprises and cloud providers to build consistent, highly secure and scalable multi-cloud networks. Already proven in the most demanding public cloud infrastructures, CloudEOS Router VM (CloudEOS-VR) extends the Arista EOS platform from Arista's physical switching and routing platforms into the virtualized cloud environment.

Arista CloudEOS-VR is available in the major public cloud marketplaces for hourly or BYOL consumption (AWS, Azure, GCP). CloudEOS-VR can also be deployed in the on-prem private cloud on standard x86 virtualisation platforms or pre-packaged on an Arista Physical Appliance Platform.

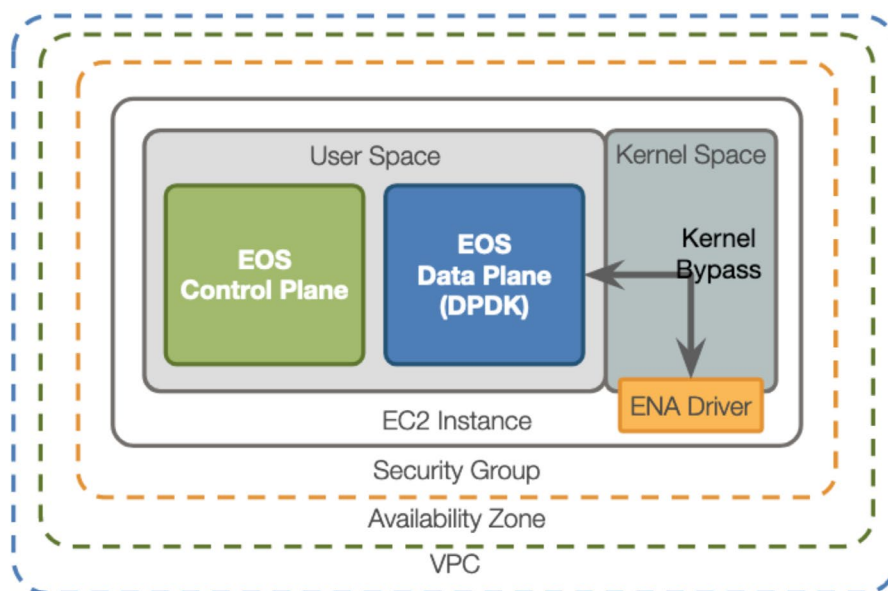


Figure 3: CloudEOS Router VM Software Architecture

CloudEOS is based on the same network operating system already proven in the most demanding public clouds, government and enterprise infrastructures, and utilizes the exact same binary image and release trains as all Arista EOS platforms, physical or virtual. Arista CloudEOS extends the Arista EOS platform with a powerful and elastic automated deployment model. This approach ensures that the CloudEOS platform will always support the latest EOS features, with the same high quality and platform compatibility as the entire Arista networking portfolio

At its core, Arista EOS provides an extremely robust, stable and resilient network-operating environment for the cloud while delivering on the need for openness, software modularity and extensibility.

The following are key attributes of CloudEOS, optimised for the multi-cloud use-cases:

- Proven Routing
- Dynamic Path Selection
- Secure Tunneling
- High Availability
- High Performance (DPDK)
- State Streaming Telemetry
- APIs and Programmability
- Multi-hypervisor and cloud-native packages
- CloudVision Integration

In a Cloud deployment context, CloudEOS may separate design patterns into three layers:

- CloudLeaf,
- CloudSpine
- CloudEdge




Figure 4: CloudEOS Design Patterns

Each layer has a set of CloudEOS Routers, enabling horizontal network scale-out, while also providing network visibility via realtime state-streaming telemetry, flow visibility and consistent operations across DC, campus and cloud. The design pattern is highly flexible, and fully interoperates with native cloud networking services like AWS Transit GW (TGW).

This model aligns well with Arista UCN (Universal Cloud Network) design for both data centers, campus and now public clouds and multi-cloud.

Depending on the customer use-cases and scale required, CloudEdge and CloudSpine can be collapsed into a single layer which will cover services such as Firewalls, IPS, Load Balancers, connectivity to other application layer VPCs / VNETs and WAN connectivity.

Using the CloudEOS Routers the network engineering team can build a highly scalable and repeatable architecture based on open standards protocols across the multi-cloud environment.

Underlay network in the public cloud

The first step in building the Arista CloudEOS architecture is deploying an underlay which allows establishing the initial router to router reachability, just like the VXLAN VTEP loopback reachability is established in an on-premises environment.

VPC and VNET peering are leveraged as the underlay network. AWS TGW can also be used for scale reasons, or any new underlay services built by the cloud provider in future. The whole underlay network bring up is part of the Terraform provisioning, which is completely automated, so you don't have to do it manually.

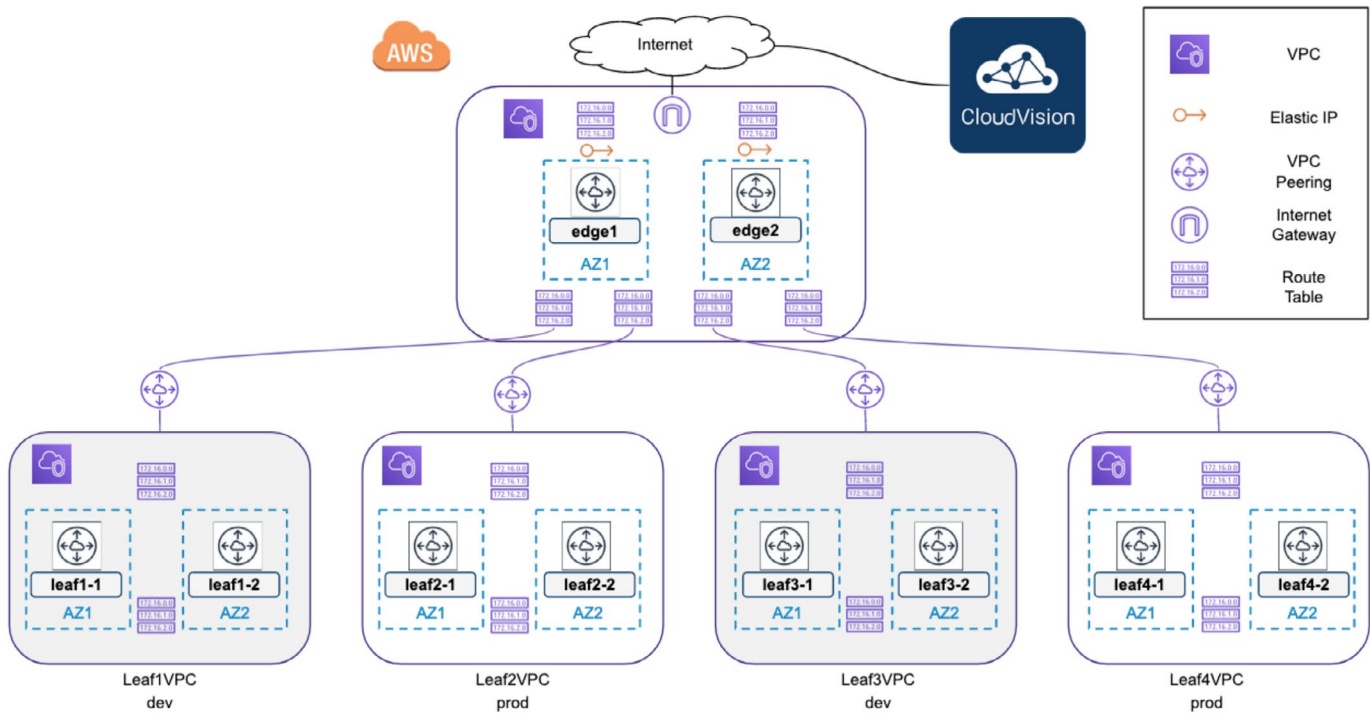


Figure 5: Typical Underlay deployment in the Public Cloud

For redundancy, the CloudEOS Routers can be deployed across availability zones. In the eventuality when an availability zone suffers connectivity issues, the Arista Cloud HA feature makes use of a BFD peering session between the two CloudEOS routers to determine connectivity loss. When this occurs, the Arista Cloud HA feature leverages the cloud native APIs to automatically remap the route tables within the failed availability zone to the remaining CloudEOS router in a different availability zone.

Overlay network in the public cloud

Once the underlay VTEP reachability is established between the CloudEOS routers, an overlay network will be built to support the creation of logical topologies in the public cloud. VXLAN is used in the dataplane, BGP eVPN as the control plane, much in the same way you would build a VXLAN fabric in your DC.

Segments such as 'dev' and 'prod' can be placed into their own dedicated VRF and can be carried cross clouds and back to the DC to maintain an end-to-end consistent segmentation.

Since these routers are just like any Arista EOS device, you can perform all the fine-grained routing control like BGP prefix list, route-map, community to further optimize your routing path.

All the overlay configurations are automatically generated by CloudVision as part of the Terraform provisioning and automatically deployed to CloudEOS Router by CloudVision, and you don't have to configure it manually. Apart from routing, the overlay VXLAN fabric allows you to provide more advanced services that cloud providers couldn't natively support or not in a consistent way, like performance monitoring, NAT for overlapping IP space issues, QoS to provide guaranteed bandwidth for mission critical applications and flow visibility.

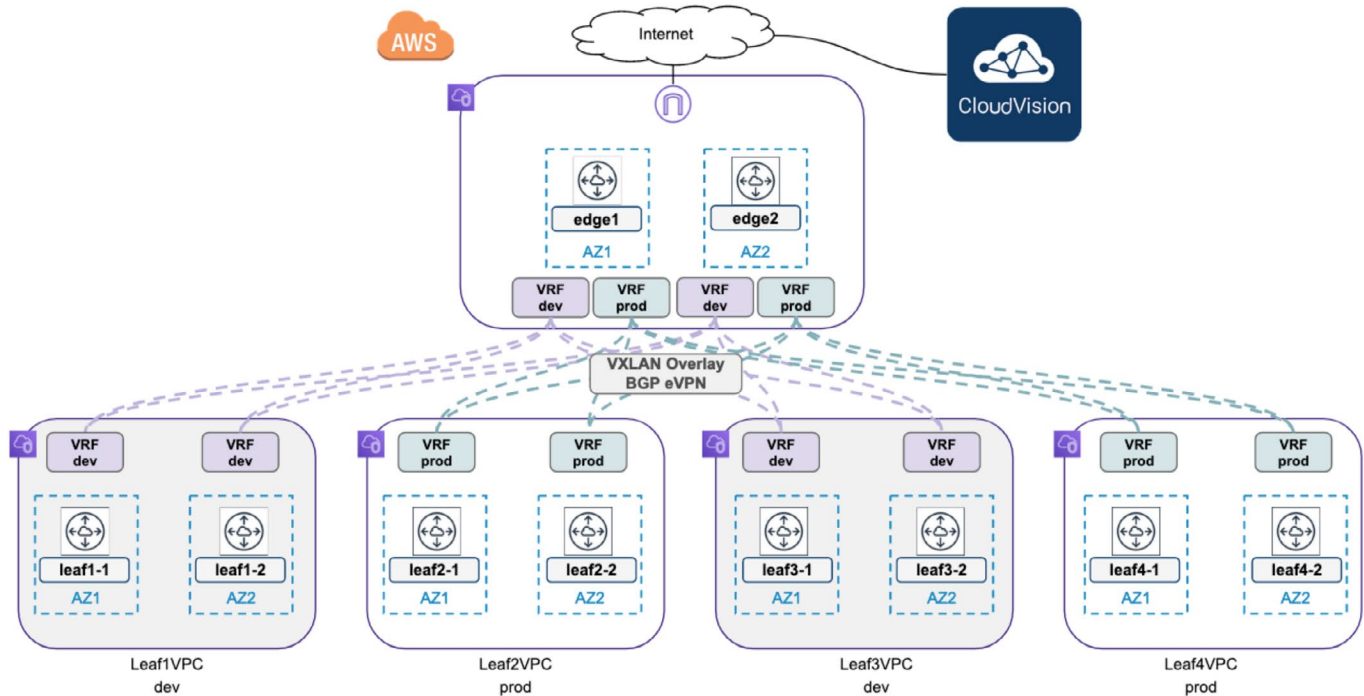


Figure 6: Typical Overlay deployment in the Public Cloud

Ultimately, extending the overlay between different public clouds and on-premises DCs will allow the creation of topologies that maintain consistency and secure segmentation across the hybrid-cloud via stretched Cloud Network Private Segments (CNPS).

Cloud Network Private Segments are large, scalable, multi-cloud VRFs glued together using the overlay and secured via IPsec encryption. Granular policies can be applied within the CNPS using application and host-aware policy while controls between segments can be enforced using high-performance virtual next-generation firewalls.

A common deployment model would have one CNPS for each Production, Staging, Development, and Test VPCs or VNET.




Figure 7: Cloud Network Private Segments

As the multi-cloud architecture is deployed, a key component of the solution is the ability to optimise and secure the cross-cloud and private DC connectivity.

To achieve this, each CloudEdge auto-discovers the available paths to the others and automatically establishes IPsec based data plane encryption. For optimised forwarding, and leveraging Dynamic Path Selection (DPS), each CloudEdge measures delay, latency, loss, and bandwidth for each potential path, and then applies this data, in real-time, to determine which path to use for which traffic class or application.




Figure 8: Secure Multi-Cloud Connectivity

2 - CloudEOS Router for Kubernetes

Arista CloudEOS Router for Kubernetes (CloudEOS-CR) provides an open and scalable solution for customers that are looking to deploy a cloud-grade routing solution for on-premises Kubernetes clusters.

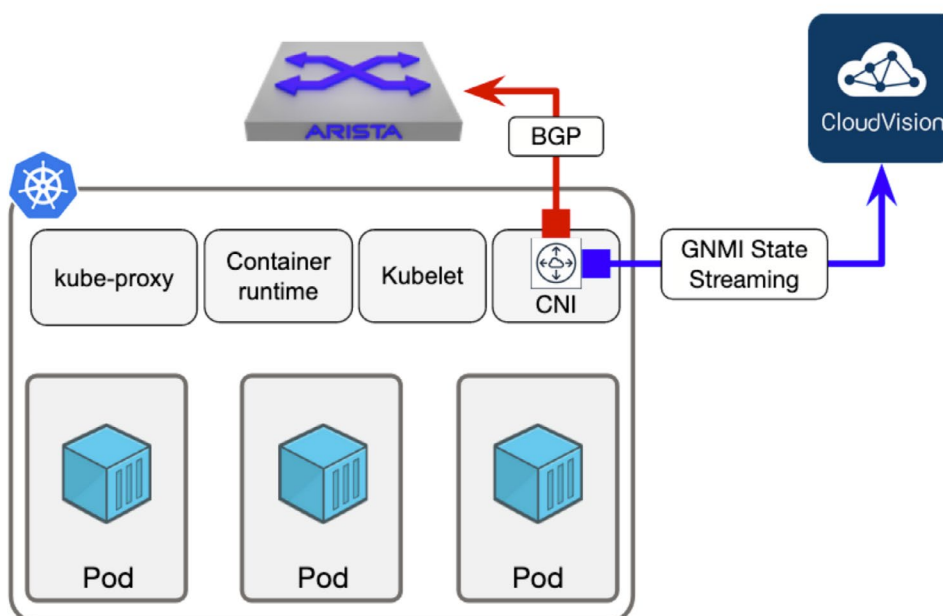


Figure 9: CloudEOS Container

To allow a congruent experience between the public cloud deployment and the on-premises Kubernetes native cloud clusters, CloudEOS-CR is based on a containerized version of Arista EOS software utilised by all Arista networking platforms. Each Kubernetes node now has access to the full power of Arista EOS and Cloudvision.

Kubernetes clusters can now be powered by enterprise-grade networking capabilities such as VXLAN, IPSEC, Packet Capture, and auto provisioning.

3 - Hashicorp Terraform Provider

To support delivering environments rapidly and at scale, a new attitude towards automating the multi-cloud solution is required. Arista has chosen Hashicorp Terraform to implement the CloudEOS solution using an Infrastructure-as-Code (IaC) approach.

Terraform is a tool for building, changing, and versioning infrastructure safely and efficiently. Infrastructure is described using a high-level configuration syntax which allows the creation of blueprints for provisioning while enabling resharing and reusing of the templates. In effect, similar to how code is maintained through a DevOps framework, the deployment can now be versioned and maintained as any other code.

Cloud workloads often require quick deployment and destruction of resources especially for transitive workloads that are created to solve a business requirement for a defined period of time. Terraform makes it easy to create the network infrastructure resources, manage the existing infrastructure and finally destroy the components that are no longer needed.

- aws_novpc
- aws_oneregion_multipleleaf
- aws_tworegion_clos
- aws_tworegion_cloudha
- aws_tworegion_noleaf
- azure_oneregion_multipleleaf

Arista provides a github repository that includes a set of cloud design patterns built and maintained by Arista. The templates provided are built based upon our past customer experiences and they form a solid base for various deployment scenarios. Customers can also modify the templates, or leverage Arista professional services for customization and integration.

<https://github.com/aristanetworks/CloudEOS/tree/master/terraform>

And more...

A new multi-cloud architecture can be deployed in three simple steps:

1. Select the Terraform template from the github repository
2. Customise it to match their requirements: number of VPCs, VNETs needed, what CIDR blocks to use, how the VPCs can be connected and placed in which CNPS for proper secure segmentation
3. Deploy the template.

Terraform will interface with the associated cloud provider via specific Terraform providers or APIs and will provision the cloud constructs such as VPCs, VNETs, subnets, route tables, as well as the CloudEOS Router VMs in each of those environments based on the design pattern or template selected at the beginning.

After the CloudEOS router instance boots up, Terraform initiates a secure connection to CloudVision for device onboarding, then CloudVision sends the EOS configurations to the CloudEOS routers to build out the overlay network.

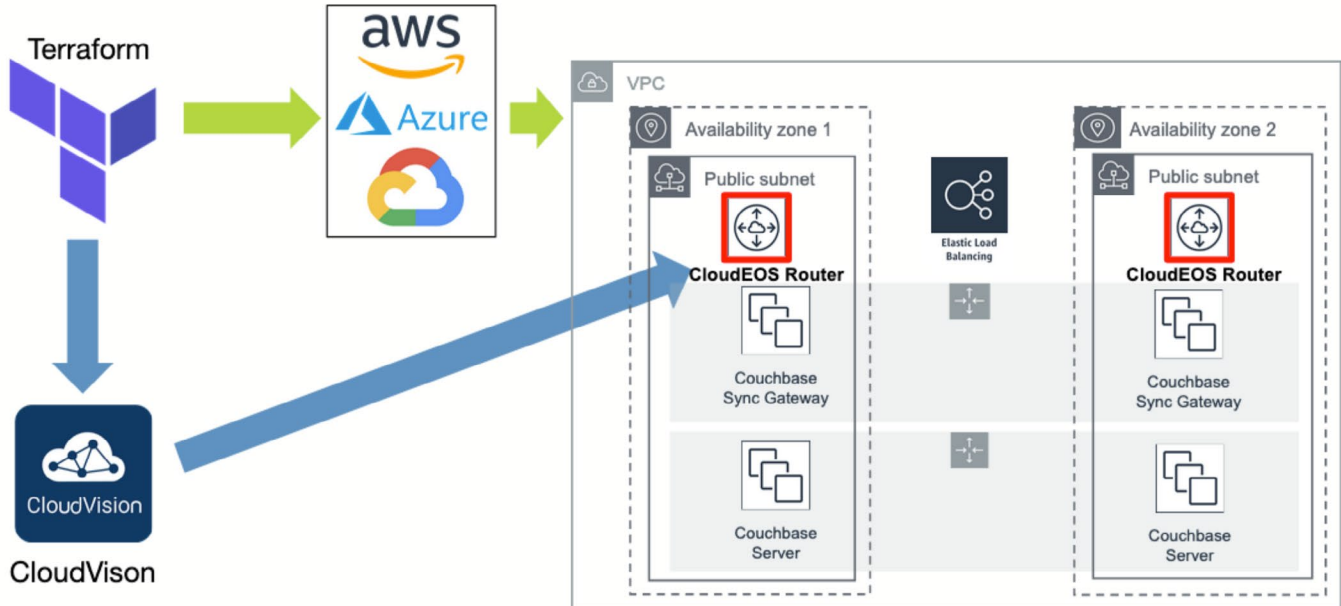


Figure 10: CloudEOS Solution Provisioning Workflow

4 - CloudVision

Arista's CloudVision is a turnkey management plane providing a modern approach to automation and telemetry. Arista Cloudvision is available as an enterprise-grade cloud-based Software-as-a-Service (SaaS) platform or available for deployment as virtual or physical appliances on-premise.

By delivering CloudVision as a cloud-based SaaS platform, customers can now have a unified and automated deployment, provisioning, and maintenance experience with no onsite resources to set up and manage. Further, Arista provides all the ongoing maintenance and tuning of the service, delivers data encryption that is always on, provides proactive security patching, and enables elastic scaling, automated backup, failover, and recovery so that customers no longer need to worry about the reliability, performance, and security of their management software infrastructure.

Cloudvision is a software product for managing any EOS instance - CloudEOS router VMs, CloudEOS routers in the Kubernetes environment, EOS running on physical Arista switches for both DC and Campus use-cases. Unique across the industry, CloudVision becomes the single management plane across all enterprise use-cases and provides correlated visibility for the data center, hybrid cloud, and even campus, helping to break down traditional box-based network silo's.

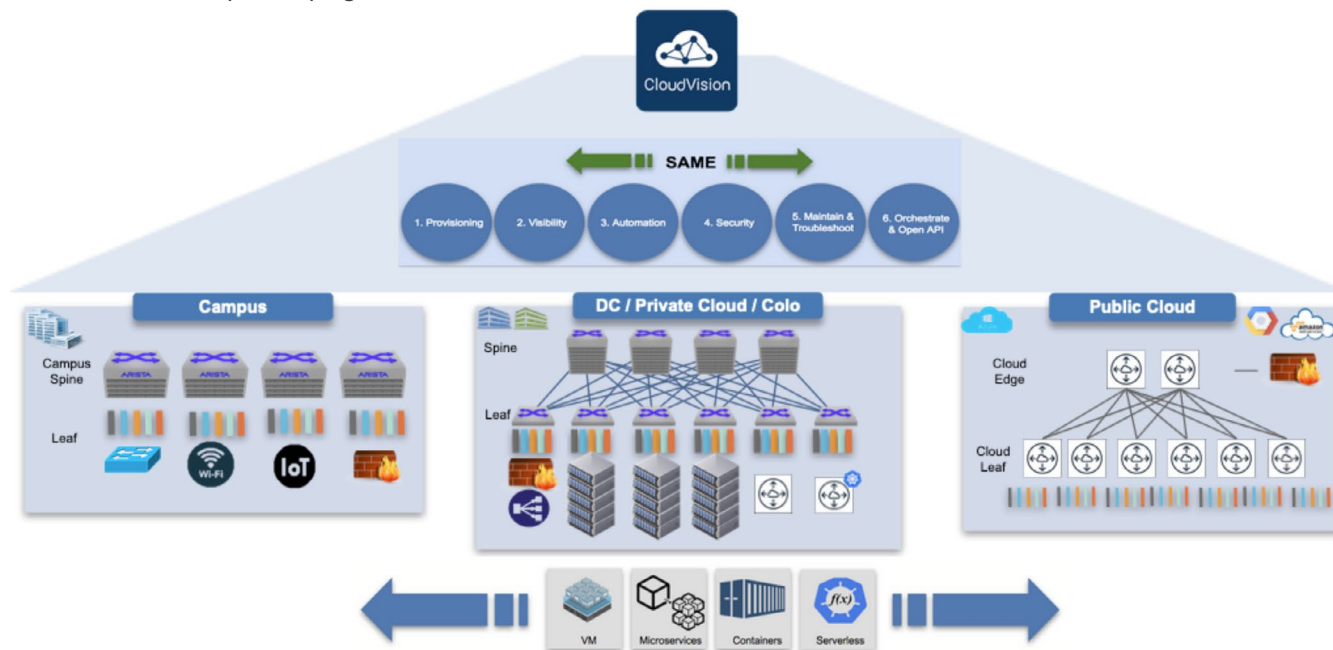


Figure 11: CloudVision - a single tool for all enterprise use-cases

As CloudEOS solution has at its core the same software image used on all Arista platforms, the same flexibility is achieved for the multi-cloud solution through using a centralized network database, NetDB, that leverages real-time state-streaming to collect an aggregate view of the entire network state. With NetDB, CloudVision becomes the point of abstraction enabling enterprise-grade network-wide automation, time-series visibility with state streaming analytics, and 3rd party integrations across the CloudEOS solution.

Cloudvision provides the following features and benefits:

- **State Streaming Telemetry:** real time streaming telemetry providing a correlated view across the multi-cloud solution, coupled with historical state for forensics troubleshooting
- **Automated Provisioning:** allows ongoing automated configuration deployments. Simple to use 'Configlets' provide modularity and simple re-use across the multi-cloud solution
- **Change Control:** automated upgrades, network rollback, and network snapshots
- **Compliance:** automatic reporting for security, audit and patch management with dashboards providing a real-time assessment of exposure to known software defects and PSIRT issues that affect the install base, and remediation recommendations
- **Open API:** RESTful APIs for all CloudVision functionality that can be used for scripting as well as integration with other management platforms and workflow tools.

Putting it all together

When deploying the Arista CloudEOS solution there are 4 simple steps involved:

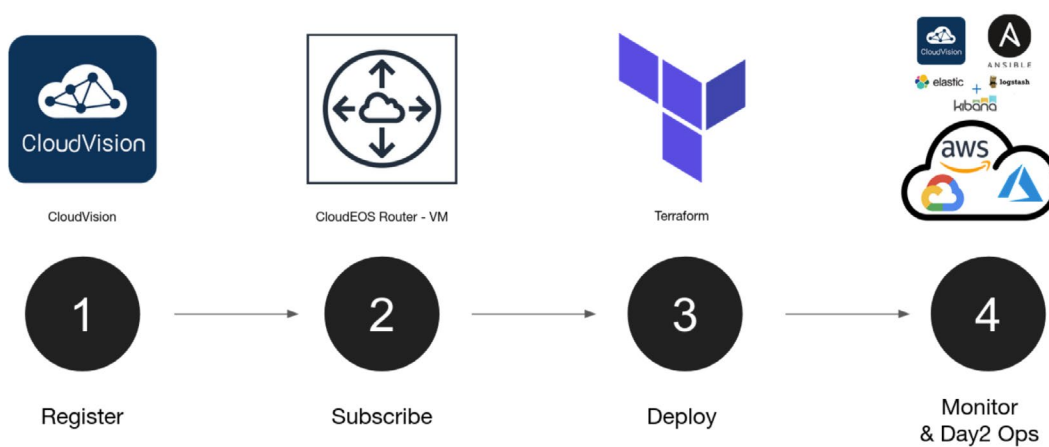



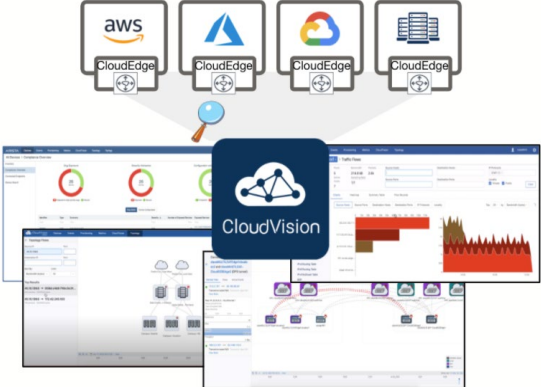
Figure 12: Steps To Build Out Your Multi-Cloud Network

1. Register for a CVaaS account on arista.io
2. Subscribe to CloudEOS Router on AWS, Azure or GCP marketplaces
3. Download Terraform Templates, customise to match your requirements and deploy the multi-cloud environment
4. Manage the lifecycle of the deployment using Cloudvision for visibility and automation

CloudEOS Use-Cases

This section provides a high level overview of the various technical use cases that are provided by the Arista CloudEOS multi-cloud solution. An organisation may deploy all or some of these offerings depending on their requirements. As such the Arista solution provides for a lot of flexibility and one can pick and choose based on their particular situation.

<p>The diagram illustrates a secure hybrid multi-cloud edge architecture. It features four CloudEdge nodes arranged in a square. The top-left node is connected to AWS, the top-right to Equinix, and the bottom-right to Google Cloud. The bottom-left node is also connected to AWS. All nodes are interconnected with secure connections, represented by padlock icons. A central cloud icon is also present.</p>	<h3>Secure Hybrid Multi-Cloud Edge</h3> <ul style="list-style-type: none"> • Secure on-demand any-to-any connectivity • Reduce latency • Direct connection between clouds without need for backhaul to the on premise DC • Robust overlay with consistent network troubleshooting toolkit
<p>The diagram shows multi-cloud path optimization. It features two CloudEdge nodes, one connected to AWS and the other to Google Cloud. A central 'Prod Traffic' box is shown. Path #1 (grey) has a latency of 30 ms, and Path #2 (green) has a latency of 50 ms. A 'Switchover' arrow indicates traffic moving from Path #1 to Path #2. Below the paths, a latency change is shown: 30 ms (green arrow) to 80 ms (red arrow).</p>	<h3>Multi-Cloud Path Optimization</h3> <ul style="list-style-type: none"> • Dynamic Path selection based on changing network conditions • Prioritize Production traffic over non-critical traffic • Complete user control allowing you to architect policies to suit your organisational requirements.
<p>The diagram illustrates consistent segmentation with central policy enforcement. On the left, a stack of 'aws' and 'Azure' logos is shown with a 'Scale out' arrow. In the center, three 'CNPS' (Central Network Policy Set) boxes are labeled 'CNPS #1', 'CNPS #2', and 'CNPS #3'. On the right, a stack of 'aws', 'Azure', and 'Google Cloud' logos is shown. Dashed lines connect the CNPS boxes to the cloud logos, indicating policy enforcement across all environments.</p>	<h3>Consistent Segmentation with Central Policy Enforcement</h3> <ul style="list-style-type: none"> • Consistent segmentation for both on-premise as well as cloud based workloads • Segmentation in the overlay using EVPN control plane • VXLAN used as the data plane mechanism • 3rd party Firewall insertion for inter-zone traffic inspection

	<h3>CloudEOS and AWS TGW Integration</h3> <ul style="list-style-type: none"> Automated deployment and provisioning with Terraform and CloudVision as-a-Service Extend TGW segmentation to multi-cloud Monitor TGW in CloudVision Integrate security services at the edge
	<h3>Visibility and Governance</h3> <ul style="list-style-type: none"> End-to-end visibility Multi-Cloud Dashboard Real time streaming telemetry Configuration management Compliance view

Secure Multi-Cloud Edge

The most common multi-cloud deployment model that utilizes cloud native tools is to backhaul traffic from the cloud to the on-premise datacenter to interconnect the customer instances in the various cloud services providers. As such the on-premise datacenter becomes the “transit” transport between the various cloud providers, leading to the need to provision bandwidth just for the transit use case, causing an increase in latency, cost and capacity planning overheads.

Using the design patterns that were introduced in the previous sections we can therefore address the primary use-case with the CloudEOS virtual appliance i.e. providing for dynamic secure, on-demand multi-cloud edge connectivity.

Arista introduces the CloudEdge CloudEOS instance to provide for dynamic secure connectivity between the various CloudEdge instances in the multi-cloud setup. We can extend this CloudEdge concept to the on-premise data center to provide for a secure hybrid multi-cloud edge utilizing dynamic IPSEC tunnels. These tunnels can be over the internet or dedicated connections such as Direct Connect or Express Route. For on-premise deployments, CloudEOS can also use paths provided by MPLS, LTE etc. in addition to the internet.

There are two main networking constructs to interconnect the various workload VPC’s and the on-premise connectivity requirements. These are achieved using:

1. **Transit VPC/VNET** to interconnect the various workload VPC/VNET’s in a certain cloud for a region. It would be recommended to use CloudLeaf and CloudSpine to form a leaf/spine architecture within a region for that specific CSP. There are two further options when it comes to deploying the CloudEdge instances:
 - a. Collapse the CloudEdge functionality into CloudSpine. This has been discussed earlier in the document
 - b. Deploy the CloudEdge within a dedicated VPC, and will act as the edge router that connects to other CloudEdge devices in other regions /clouds. When you compare this to the on-premise UCN architecture, this dedicated VPC is equivalent to the border/service leaf switch.

Most customers will end deploying option 1a i.e. collapse the CloudEdge into the CloudSpine as this will prove to be a more cost effective option. In either case, the connectivity between the various clouds will be between the CloudEdge instances.

2. **AWS Transit Gateway (TGW)** acts as a central hub and interconnects workload VPC's and will also provide for connectivity to the on-premise deployment. Arista provides for integration with TGW and we have provided for more details later in this document in the use-cases section.

The various CloudEdge instances in the various clouds need to autodiscover each other to form the BGP EVPN control plane and then subsequently enable the VXLAN for data plane purposes

In order to keep the solution simple and open-standards based, Arista introduces the concept of a route reflector in the design. The route reflector is a standard BGP implementation which reduces the requirement for a full mesh of BGP EVPN sessions between the CloudEdges.

As such the steps for creating this solution are as follows:

1. Deploy the CloudEdges using the into the various clouds
2. Deploy BGP route reflectors (RR) for scaling the solution. These are referred to as cloud RR.
3. The CloudRR's can be deployed on premise or in the cloud. There is complete freedom in choosing where you deploy them. For redundancy we would recommend deploying at least 2 RR's
4. To create the sessions between the CloudEdge and the CloudRR, we then need to use Dynamic Path Selection (DPS) between the CloudEdge and the CloudRR. These DPS paths can be secured using IPSEC tunnels if so required
5. Once these are setup we can overlay EVPN sessions between the CloudEdge and CloudRR's
6. The CloudRR then discovers the various CloudEdges that make up the setup and allow the CloudEdges to create secure VXLAN based tunnels for the payload directly between themselves.

All of the above provisioned declaratively using Terraform provider templates and CloudVision in concert to provide an automated deployment without any manual intervention.

There are a host of technical benefits using this approach. Some of these are listed below:

- Reduction in the number of BGP EVPN sessions. Each CloudEdge has as many sessions as there are CloudRR instances
- On account of DPS there are multiple paths between the CloudEdges and the CloudRR providing for resiliency for the control plane
- The dataplane VXLAN tunnels are formed directly between the CloudEdge instances and the CloudRR is never in the datapath.
- The VXLAN tunnels can also be optionally encrypted with IPSEC and as such provides for security and privacy.
- Using the EVPN control plane we can overlay multiple CNPS between the various clouds.
- Since this deployment is an example of L3 eVPN, type 5 routes are advertised by each site which allows the routing tables to learn of the other subnets/CIDR blocks in the various CNPS.

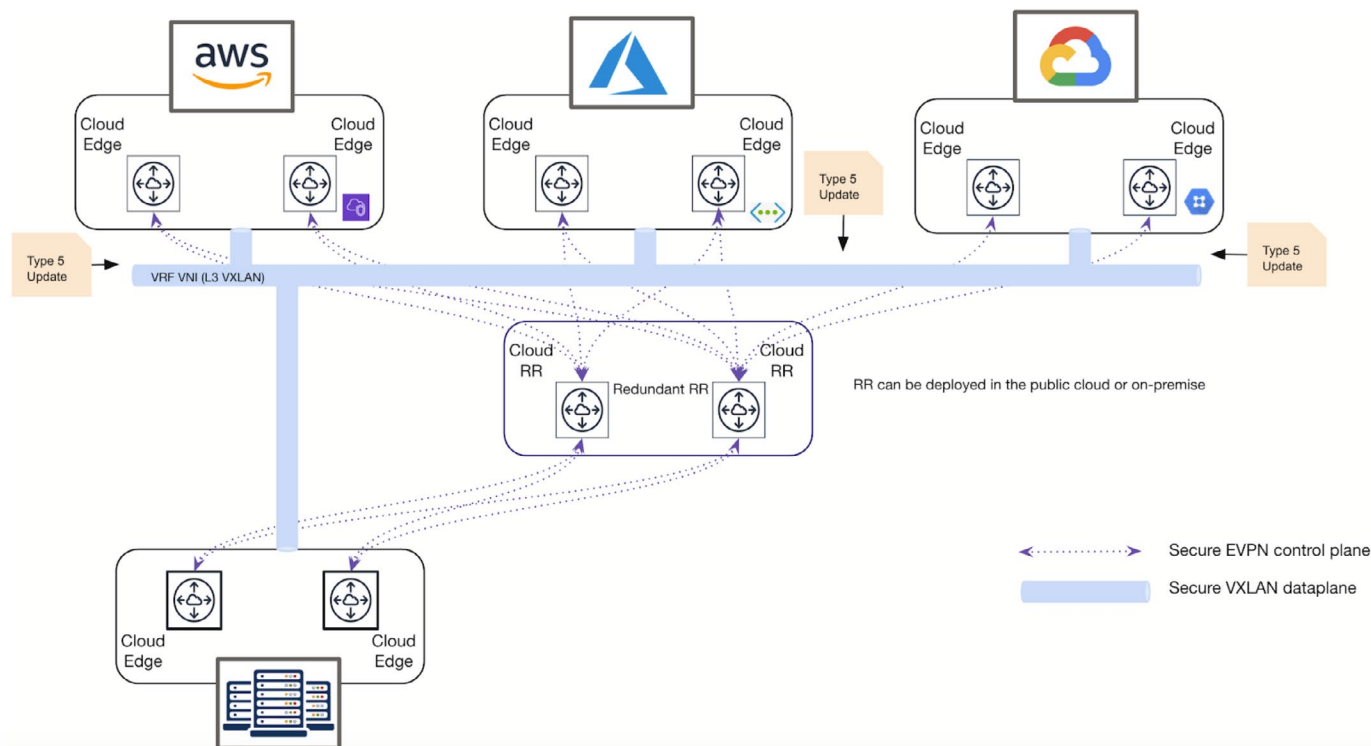


Figure 13: CloudEOS Fabric Control Plane and Data Plane

The data plane used in this architecture is VXLAN. The CloudEdge instance as such acts as a Layer 3 VTEP with the VNI providing for VPN-capabilities and acts as a CNPS/VRF lookup key to ensure that routing lookup happens in the correct CNPS/VRF. This use-case is discussed later in this "segmentation" use case.

In order to keep the solution simple and scalable without operator overhead, Arista further utilizes BGP to dynamically discover sites along with IPSEC key generation and exchange. The ability to dynamically discover sites forms the basis of the multi-cloud path optimization use-case discussed as part of the "multi-path optimization" use case.

Arista secure multi-cloud solution provides for the following benefits:

- Normalize the network connectivity across the multi-cloud landscape
- Provide for repeatable, scalable, consistent design patterns
- Declarative provisioning via Terraform Hashicorp
- Dynamic routing updates and a rich set of network troubleshooting tools
- Consistent API's across the environment
- Consistent operational experience leading to operational efficiencies
- Streaming telemetry to unify the information from the various CloudEOS instances

Multi-Cloud Path Optimization

Multi-cloud path optimization refers to the ability of identifying and dynamically selecting the best path(s) between the various CloudEdges. This use case is therefore the natural progression to hybrid multi-cloud connectivity discussed earlier.

The basic concept of Dynamic Path Selection (DPS) is to select a path or paths between CloudEdge instances connected via secure IPSEC tunnels. These secure tunnels carry VXLAN encapsulated packets based on the control plane information provided by BGP EVPN.

The DPS Tunnel is therefore really an abstraction of all possible paths between the two CloudEdge nodes where the actual path chosen is a function policy and path telemetry information.

The figure shows that there exists multiple candidate paths between the AWS instance and the other clouds. DPS will select the path(s) based on application policy and will send traffic over the selected paths. As such we can carve out different links or paths based on the applications or even the different environments in a customer setup. Path characteristics that can be specified include latency, jitter, delay, packet loss and available bandwidth.

This figure below is an overly simplified view to get the point across. Also note that the diagram below only shows the DPS paths for the IPSEC tunnels for the VXLAN overlay. A similar construct will be present for DPS paths to the CloudRR for EVPN peering.

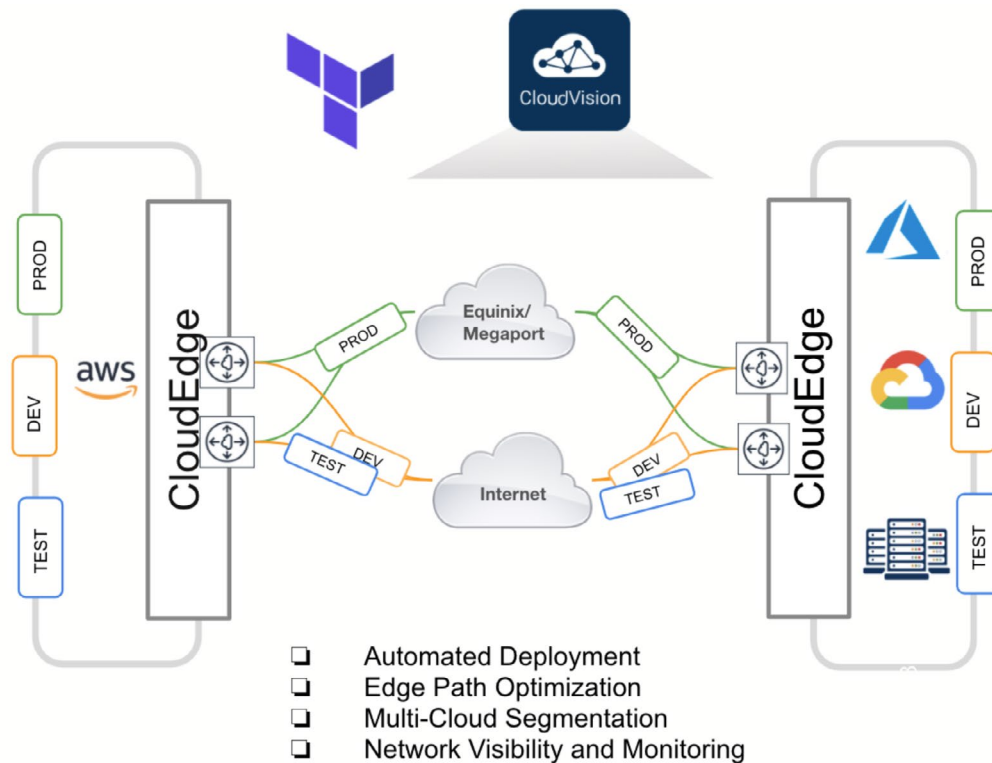


Figure 14: Multi-Cloud Connectivity with Multiple Available Paths

The combination of application profile and dynamic paths allows the operator to select paths based on the requirements that are unique to an organization. IPsec can be then added to those paths in the event that one would like to encrypt these sessions. In the internet use case shown above, IPsec becomes mandatory, however for dedicated circuits one may want to keep IPsec as optional and therefore we have provided the option to the customers to pick and choose features which make sense to them.

If multiple paths meet the criteria then the traffic is load balanced using equal cost multipathing (ECMP). Path preferences can be specified to select one path over another. As such the DPS feature allows one to carve out policies that ensure optimal performance for important applications without degradation to meet SLA targets in the face of changing network conditions without user intervention. In addition, one can leverage QoS, NAT etc if required to further optimize that traffic between the various clouds for a truly enterprise grade solution.

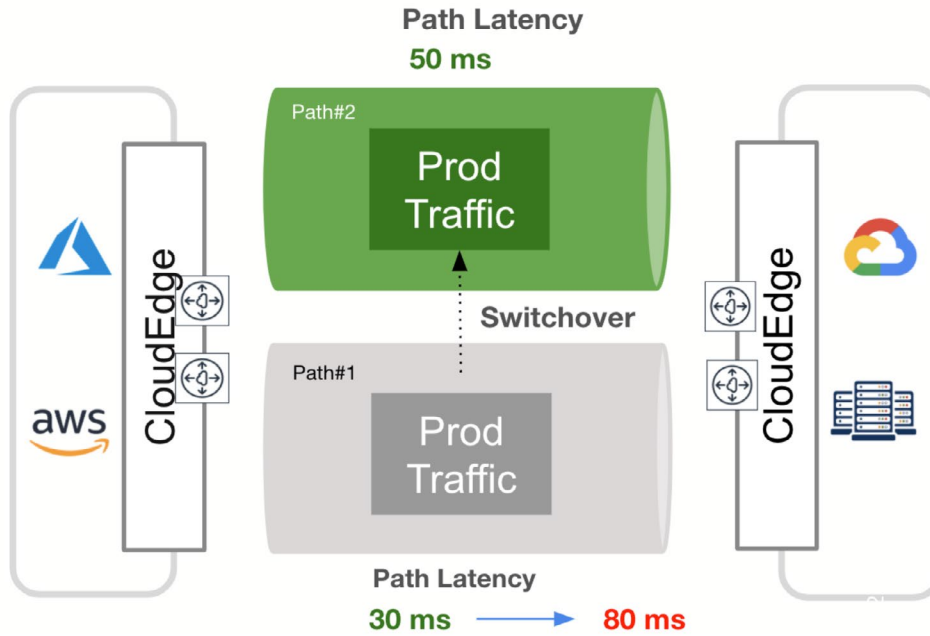


Figure 15: Multi-Cloud Path Optimization For Business Critical Applications

Consistent Segmentation with Central Policy Enforcement

Arista’s strategy for network segmentation in both private and public clouds is based on open standards protocols and repeatable design patterns. A typical private cloud’s network segmentation strategy may look something like this:

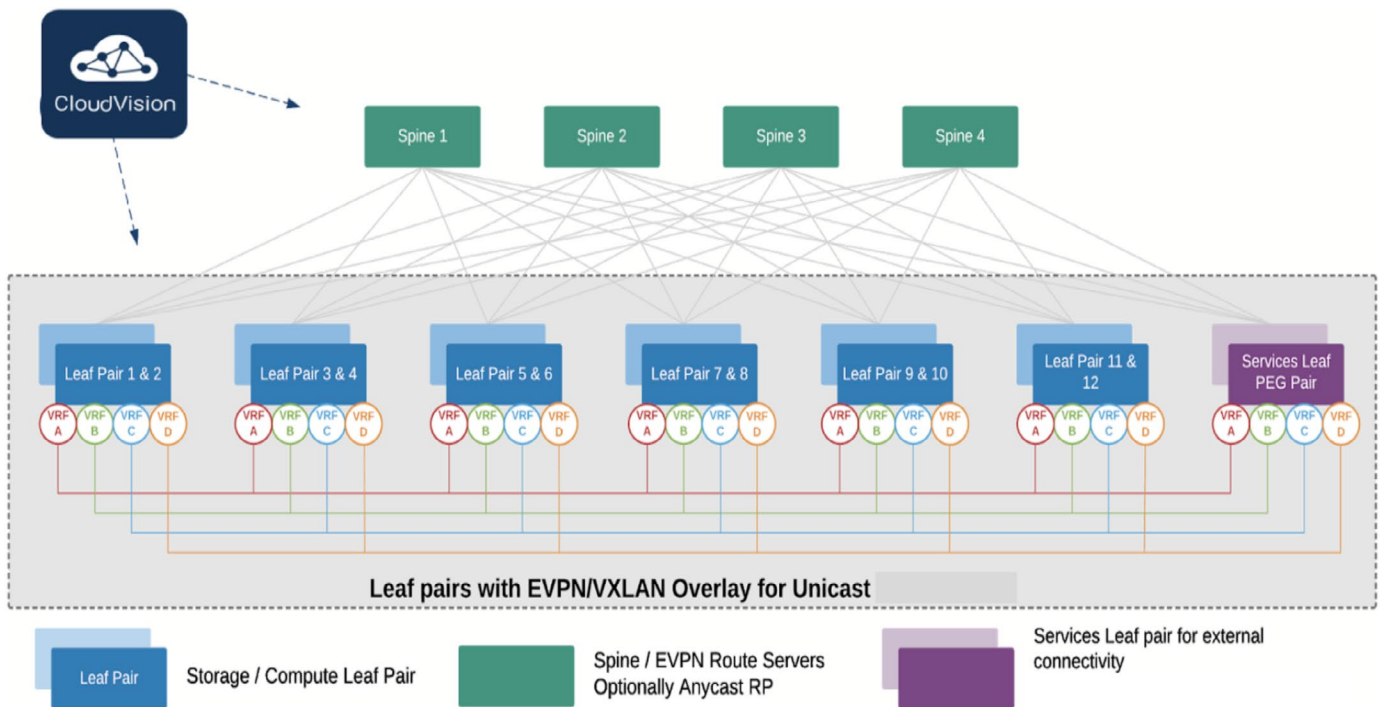


Figure 16: Network Segmentation in Data Center Using EVPN/VXLAN

The spine layer provides connectivity between the leaf layer and different leaf layer subnets are mapped to different VRFs (aka Virtual routing and forwarding). Inter-VRF as well as External connectivity happens at the Services leaf layer. A public cloud implementation can employ a similar strategy managed by Arista’s Cloudvision Portal.

This strategy includes network-wide segmentation with a VXLAN (Virtual Extensible LAN) overlay and Border Gateway Protocol with eVPN as the NRLI (Network Layer Reachability Information). This open standards based implementation using BGP allows the customers to segment and inspect traffic with a firewall vendor of their choice.

Some of the benefits of using eVPN

- It is underpinned by BGP, which is open-standards based, most commonly used in most cloud infrastructures.
- It supports Layer-3 mobility services. (note that public clouds suppress Layer-2 traffic)
- It removes complexity and risk from critical aggregation points.

Given that we encourage employing VXLAN with BGP eVPN and all public clouds support VXLAN encapsulation, L3VPN services can very easily be extended between multiple clouds.

Multi-Cloud Network Segmentation

As part of the CloudEOS fabric bring up, which is based upon VXLAN and BGP eVPN, enterprise customers can map their cloud resources like VPCs, or VNETs, or subnets into different VRFs for network segmentation across multiple cloud providers. For example, DEV VPC in AWS can only talk to DEV VNET in Azure, but not the PROD VPC or VNET. For inter-segments communication that requires inspection, CloudEOS can route that traffic to a Central Policy Enforcement VPC, or a VNET, which may leverage a set of firewalls for traffic inspection purposes. With different design options, we can scale out the firewall clusters based on the performance needs. Inter-segments communication that doesn't need inspection, like shared services such as logging or active-directory that needs to be accessed from all segments, can be enabled through the VXLAN fabric via route leaking across different VRFs.

As described in the previous section related to how the VXLAN fabric is built in the public cloud, the underlay network can be chosen based on the scale and performance requirements, all the routing intelligence happens in the overlay network. Enterprise customers can easily build out a consistent segmentation solution, as shown below, which uses AWS TGW (more details with TGW is explained in the next section, TGW can be replaced with VPC peering or other techniques in Azure and GCP) as the underlay network.

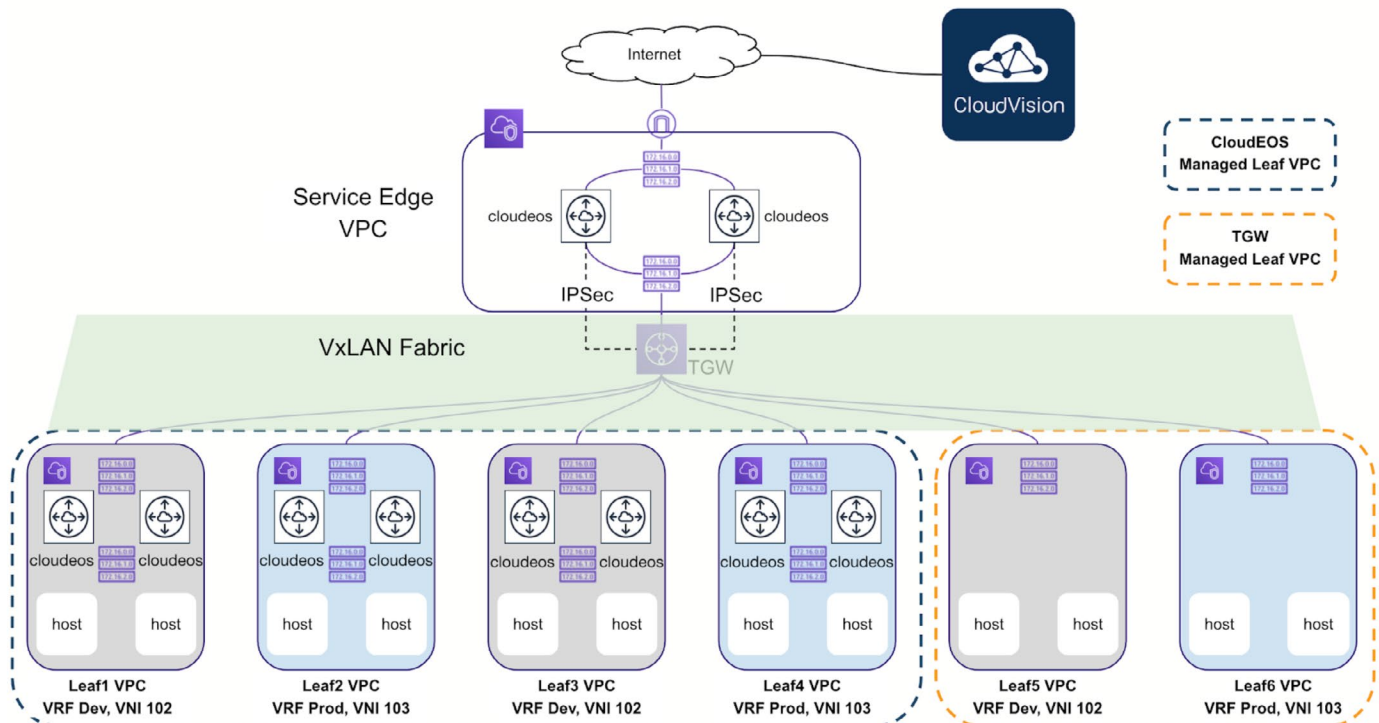


Figure 17: Network Segmentation in Public Cloud

Network Segmentation Goal

Leaf1, Leaf3, and Leaf5 VPCs belong to the DEV environment, Leaf2, Leaf4 and Leaf6 VPCs belong to the PROD environment. All Leaf VPCs 1-6 have TGW VPC attachment associated with.

CloudEOS Managed VPCs

Leafs 1-4 demonstrate consistent overlay architecture that can be deployed across any cloud provider. CloudEOS routers in Leaf 1-4 are placed into a single TGW route table that enables basic ip reachability for all the CloudEOS routers to build out the VXLAN fabric and all segmentation happens within the VXLAN overlay fabric.

TGW Managed VPCs

Leaf 5 and 6 demonstrate the native segmentation capability provided by TGW, with placement into different TGW route tables, and without any CloudEOS Routers in those VPC's.

Service Edge VPC (where Firewall resides)

Service VPC, with its VPC attachments to Leaf 1-4 and Leaf 5-6, becomes the edge point for customers to insert services like firewalls, etc.

How Network Segmentation is Implemented

From Leaf 1 to Leaf4, each Leaf is logically mapped to a unique VRF (Virtual Router Forwarding), then gets mapped to a VXLAN VNI (Virtual Network Identifier), which ensures that traffic in the overlay stays segmented. The Ethernet interfaces shown in this topology in Leaf1 and Leaf3 are configured in VRF DEV - VNI 102 and Leaf2 and Leaf4 get mapped to the PROD VRF - VNI 103.

From Leaf 5 to Leaf6, since there is no CloudEOS in these VPCs. The CloudEOS in Service VPC has VPN attachments with multiple IPSec tunnels to TGW for route exchange. Each IPSec tunnel is placed into respective VRF on CloudEOS in Service VPC. For example, IPSec tunnel to TGW route table that has Leaf5 is placed under VRF DEV - VNI 102, IPSec tunnel to TGW route table that has Leaf6 is placed under VRF PROD - VNI 103.

With that, Leaf 1,3,5 can talk to each other in the DEV environment, but not to Leaf 2,4,5 in the PROD environment.

Firewall Insertion for Central Policy Enforcement

A natural step after network segmentation is central policy enforcement. There are multiple factors to consider when designing a policy enforcement solution in the public cloud:

- What type of traffic
 - › Inter-segments, east-west/north-south communication
 - › Internet connectivity, compliance reason
- What approach
 - › Distributed vs centralized
 - › Vendor Firewalls vs Native Cloud Firewalls
 - › Active-Standby vs Active-Active
 - › Traffic symmetry (S-NAT vs no S-NAT)
- Network topologies
 - › Hub-spoke, full-mesh, transit gateway

At Arista, we provide solutions that work with all the design considerations. In the example below, we are showing one of the solutions with active / active firewalls inserted in the Services VPC, that provides enterprise customers the ability to allow / disallow Inter-CNPS or VRF traffic (for example, in our topology between the DEV and PROD segments).

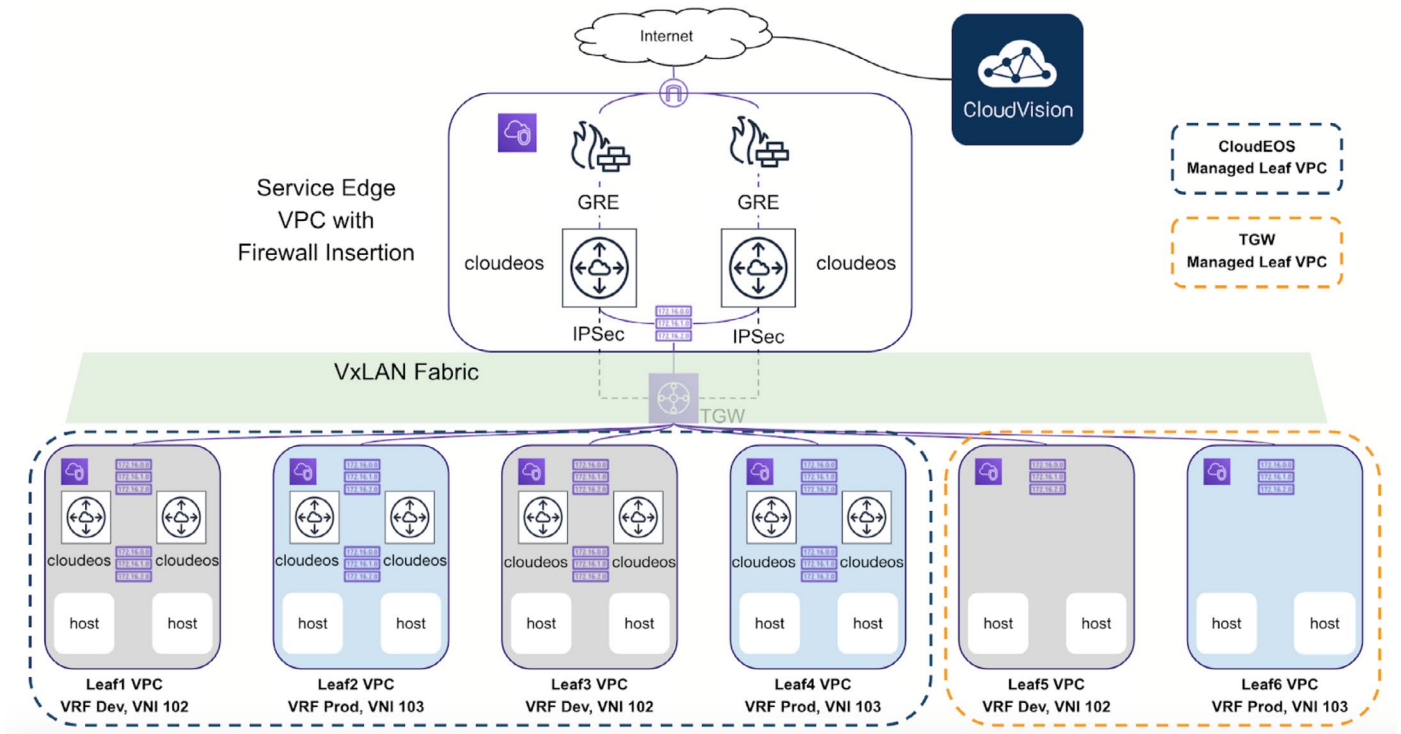


Figure 18: Firewall Insertion in Public Cloud

We expanded the previous setup to include a pair of firewalls (any vendor) in the Service VPC. Using dedicated ethernet interfaces and GRE tunnels to abstract routing from the underlay between CloudEOS Routers and the firewalls. These GRE tunnel interfaces reside in DEV and Prod VRF respectively. DEV VRF tunnel interfaces on the CloudEOS Routers advertise prefixes for Leaf-1/3/5 VPC's and PROD VRF tunnel interfaces on the CloudEOS Routers advertise prefixes for Leaf-2/4/6 VPCs. GRE Tunnel interfaces on the firewalls are all in the default VRF.

With the firewalls advertising the DEV prefixes to PROD VRF and vice versa in this topology we can control what to allow and filter traffic between the VRFs or CNPSs. This solution can also be used to provide Internet access for all Leaf VPCs.

CloudEOS and AWS Transit Gateway Integration

Transit Gateway (TGW) is a networking service provided by AWS to interconnect VPCs, data centers and remote sites. It is highly scalable, but has limited feature sets. Integrating CloudEOS with AWS TGW enables enterprise-level features, multi-cloud routing and network visibility with seamless automated provisioning and deployment for enterprise customers. Customers can choose either of the two options below or mix-and-match based on their use-cases.

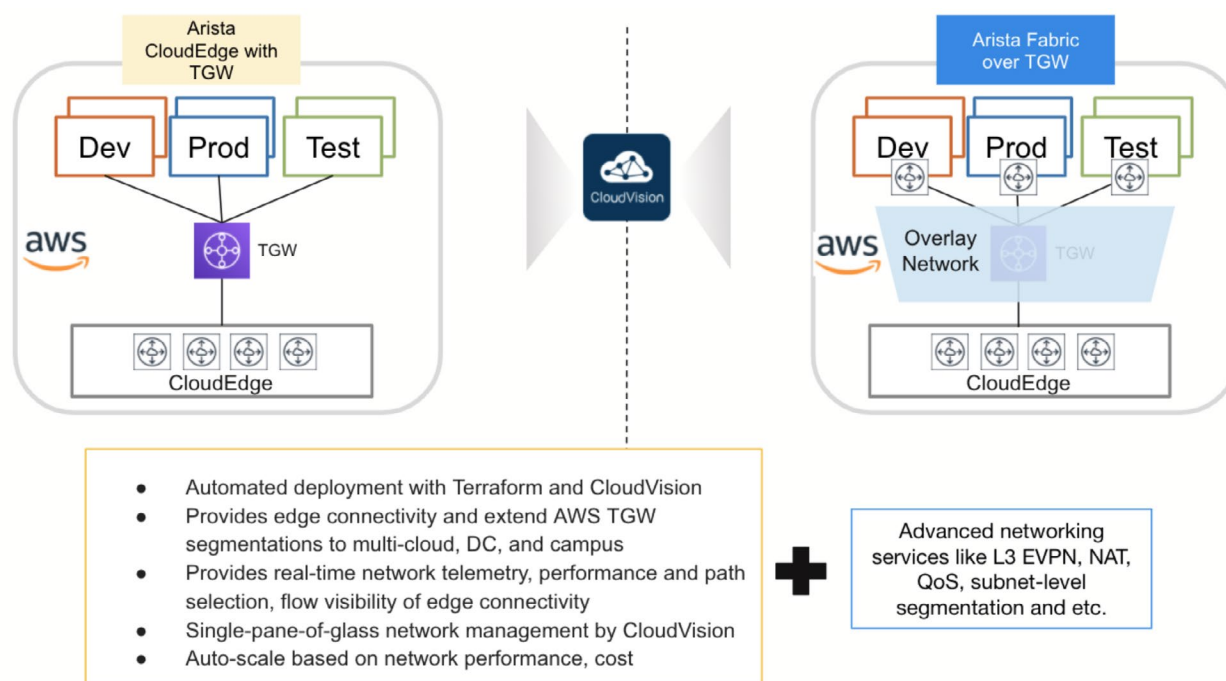


Figure 19: Integrating CloudEOS with AWS Transit Gateway (TGW)

Arista CloudEOS Edge with TGW

In the left option above, customers usually use TGW for network connectivity and segmentation between VPCs. Arista CloudEOS Edge provides edge connectivity between AWS TGW and on-prem DC, Campus and other public clouds. CloudEOS Edge router is connected to TGW using IPsec tunnel and exchange routes with BGP. For customers that are using TGW route domain for segmentation, they can extend TGW route domain to Arista CNPS (Cloud Network Private Segment) in customers' existing environments. For example, Dev VPC in AWS can only communicate with Dev VNET in Azure, and Dev resource in Arista on-prem DC, but not with Prod or Test environments unless through a central security device for inspection. Arista provides the Terraform template that automates the CloudEOS Edge deployment with TGW. Customers can monitor TGW, VPCs and CloudEOS Edge in CloudVision for network visibility, and flow visibility.

Arista Fabric over TGW

In the right option above, customers can provide more advanced enterprise services that native TGW could not support, like L3 EVPN, NAT to address overlapping ip space issues, subnet-level segmentation and QoS. CloudEOS Router will be placed as a routing gateway into customer's workload VPC which becomes CloudLeaf VPCs. CloudLeaf and CloudEdge VPC are connected using Arista DPS (Dynamic Path Selection) and BGP-EVPN on top of AWS TGW. This is a truly cloud-agnostic architecture that can be replicated into Azure and GCP easily. All the deployments are automated by Arista Terraform and CloudVision.

Mix-and-match

Depending on each application's requirements, customers can deploy "Arista CloudEOS Edge with TGW" for development and testing environments where standard AWS networking services are sufficient. For business critical applications like Prod environments, customers can deploy "Arista Fabric over TGW" to ensure applications performance, SLA and security. This is used in the previous network segmentation section that Leaf1-4 VPCs are using Arista Fabric over TGW, whereas Leaf5-6 VPCs are using Arista CloudEOS Edge with TGW.

Visibility and Governance

One of the challenges that organisations have is to collect visibility information from the multiple cloud environments that make up the hybrid multi-cloud deployment model and make sense of that data in a single dashboard that provides for an end to end view.

Every CSP provides for its own tool set to provide for a view into the networking aspects of cloud. However these are discrete data sets that need to be normalized and ingested into a customized platform making it hard from an administrative point of view.

Secondly, each cloud provides management visibility at different aggregation time periods making the job of co-relating associated flows extremely hard if not outright impossible.

What is required is a turnkey solution that provides for management of the entire setup - on premise as well EOS instances in the cloud. This is where CloudVision comes into the picture.

CloudVision provides for network management of the entire network estate including on-premise as well cloud based workloads. Some of the key features from a visibility & governance standpoint are:

- **Network Topology view:** Provides for a real-time view into the network topology and devices connected to it. Cloud Vision allows you to overlay information on top of the topology to view parameters like bandwidth utilization, throughput, errors, discards, CNPS/segmentation view etc. to get a bird's eye view into the network.

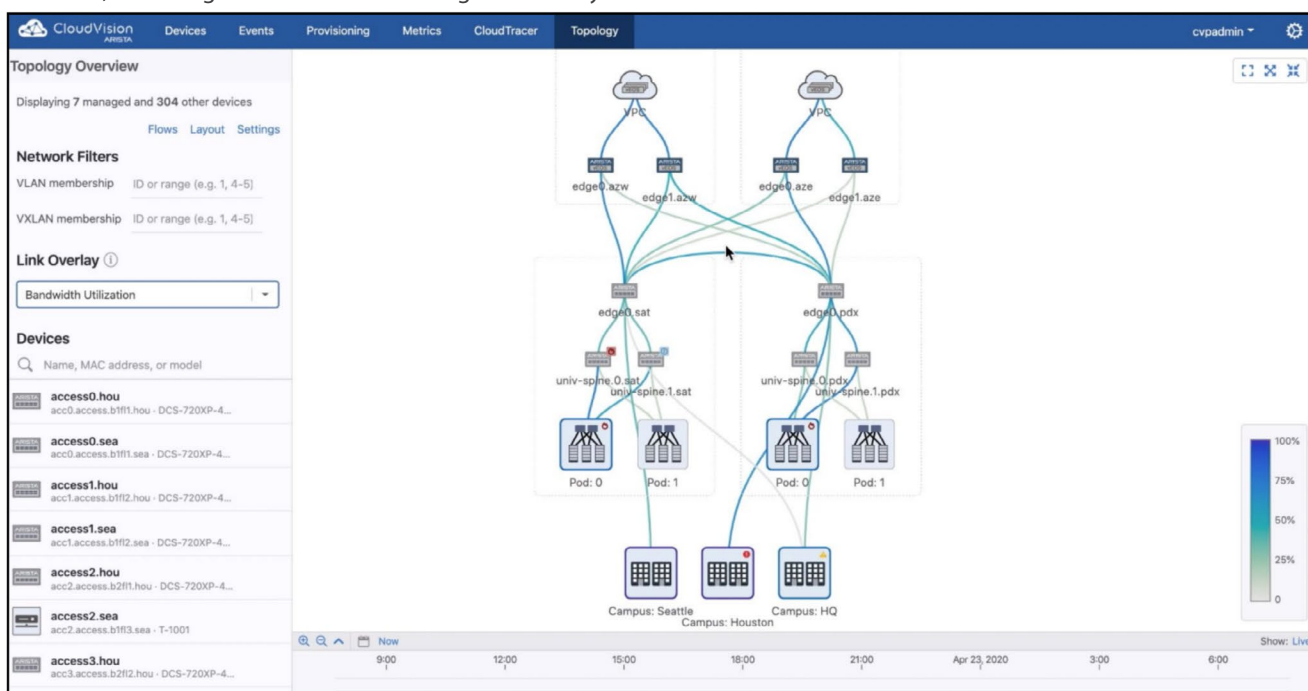


Figure 20: Network Topology View in CloudVision

- **Cloud Specific Dashboards:** This feature provides a single pane of glass which allows you to see relevant information across your multi-cloud setup. The intention behind these dashboards is to show information relevant to the networking environment for an enterprise. CVP and CloudEOS work in concert to provide a view across the various CSP's and visualize the state of the network between them. This includes:
 - › VPC/VNET details
 - › CIDR Blocks
 - › Account information
 - › CNPS details
 - › Path characteristics
 - › Connectivity details

Region ↑	VPC Name	Network Role	CIDR	Segment	CloudEOS	Account	VPC ID
us-east-1	fi-test-aws1-Leaf3Vpc	Cloud Leaf	103.2.0.0/16	dev	Yes	631918477817	vpc-03fbb1b0387bf9919
us-east-1	fi-test-aws1-EdgeVpc	Cloud Edge	100.2.0.0/16	dev, prod	Yes	631918477817	vpc-0438e8ea19ba5ee ca
us-east-1	fi-test-aws1-Leaf4Vpc	Cloud Leaf	104.2.0.0/16	prod	Yes	631918477817	vpc-03ef4c5db15c47e3c
us-east-1	fi-test-aws1-Leaf2Vpc	Cloud Leaf	102.2.0.0/16	prod	Yes	631918477817	vpc-0de747efa275611af
us-east-1	fi-test-aws1-Leaf1Vpc	Cloud Leaf	101.2.0.0/16	dev	Yes	631918477817	vpc-010052f7a8550f7f8

Figure 21: Multi-Cloud Dashboard in CloudVision

In addition to the dashboards, real-time topology views have been updated to include information on DPS and the CNPS across the entire network. This feature provides for a visual representation of the setup with the necessary information overlaid on the topology for easier verification and troubleshooting.

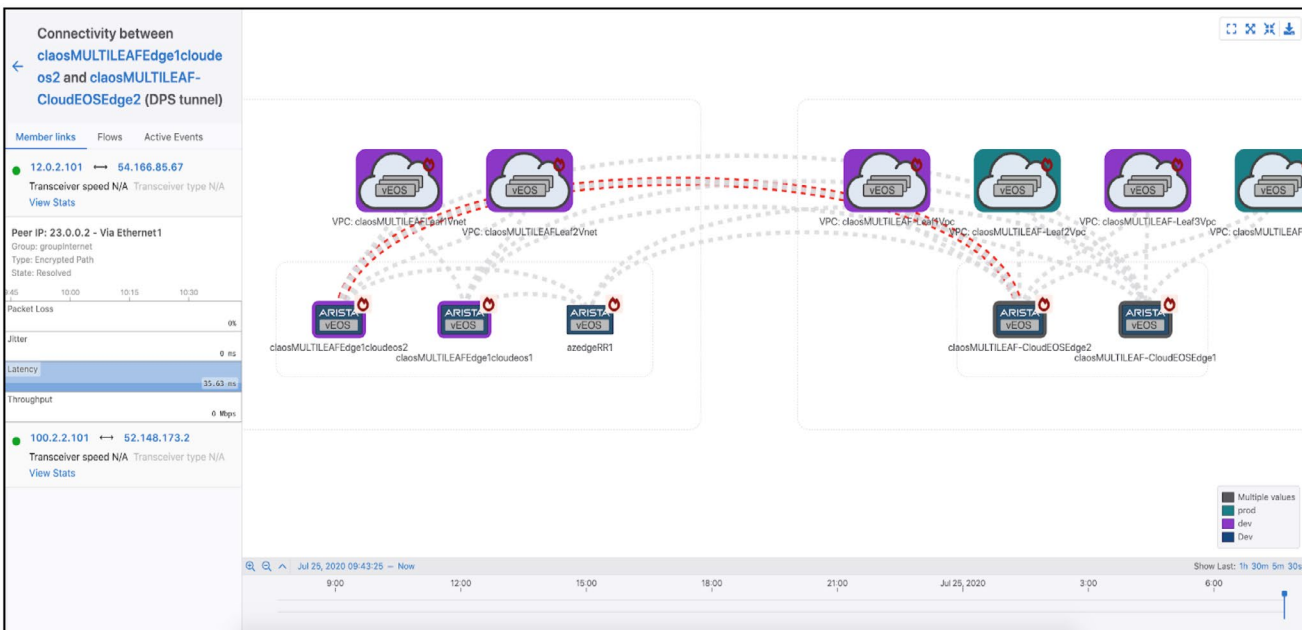


Figure 22: Topology View with Cloud Segments and DPS Path Information in CloudVision

- **Endpoint Flow:** Endpoint flow provides for an EOS specific as well as a network-wide view into the application flows - including top talkers, heat maps as well as path of that flow for an enterprise-wide view into the flows. This is enabled via IPFIX/SFlow being sent over to CVP.




Figure 23: Flow Tracking View in CloudVision

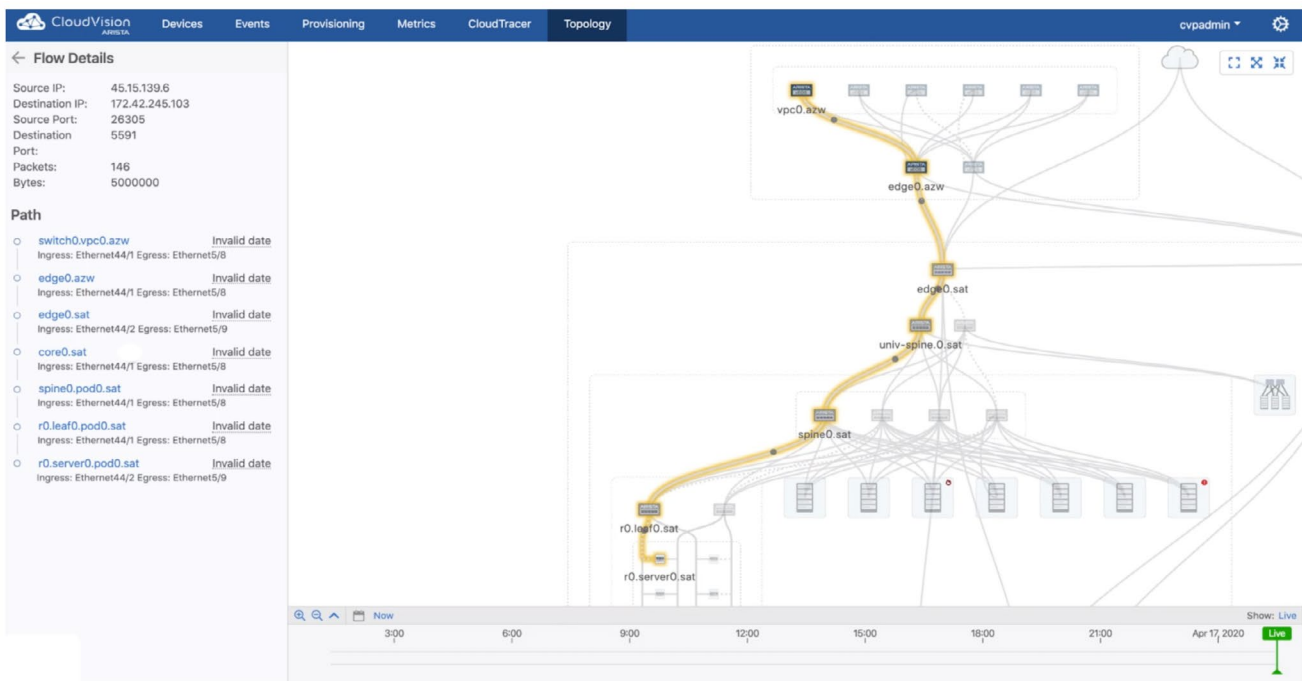


Figure 24: End to End Flow on Topology View in CloudVision

- Network-wide Event View:** CloudVision receives streaming telemetry from all the devices registered to it. The backend analytics engine compares events and categories faults based on their criticality thereby removing “noise” allowing you to concentrate on issues/events that need one’s attention.

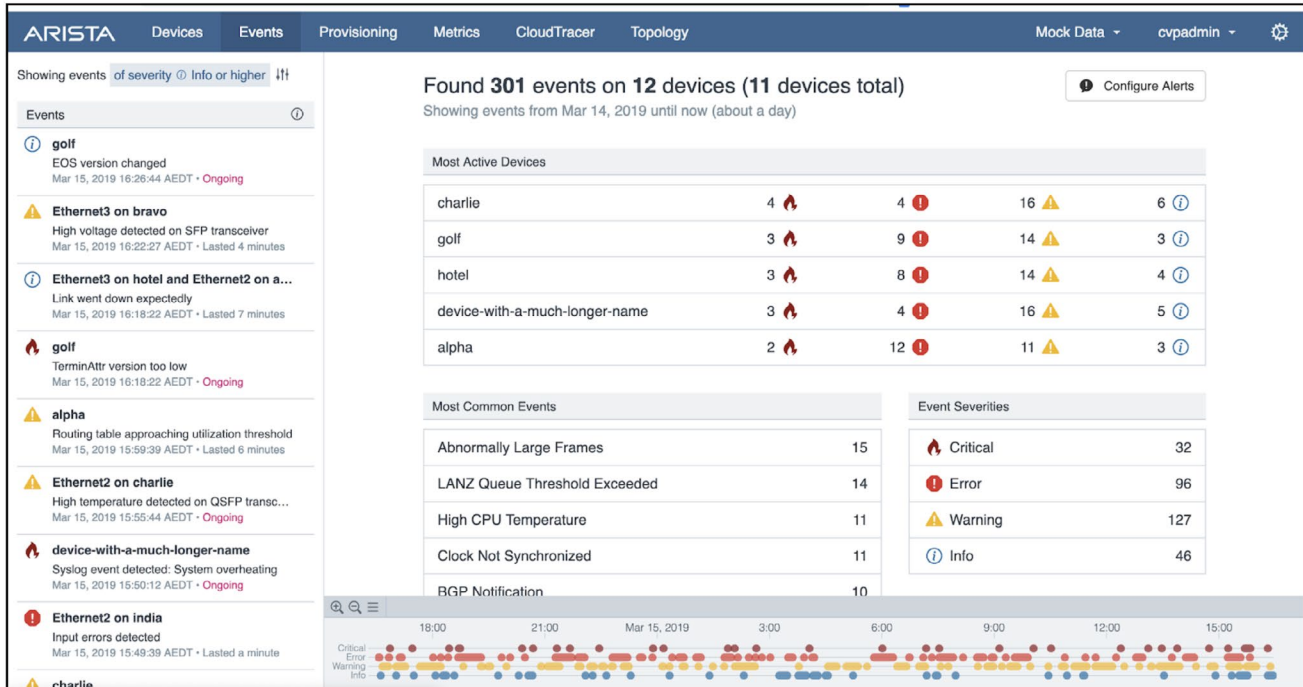


Figure 25: Events View in CloudVision

- Compliance View:** Compliance view allows you to see the compliance of your network. This includes Bug Alerts, Security vulnerability alerts, Configuration and software image compliance across all installed Arista assets in the customer's environment.



Figure 26: Compliance Dashboard in CloudVision

- Anomaly Detection:** CloudVision provides for predictive analysis where the CloudVision Analytics engine uses advanced AI/ML algorithms to proactively detect deviations from the baseline before this change causes widespread degradation in service.



Figure 27: Anomaly Detection in CloudVision

- **Address Search:** This feature allows you to search for an endpoint in the hybrid multi-cloud setup. This endpoint could reside in any cloud - be it public or on-premise private cloud.

Address Search

Search for MAC and IP Addresses here to find Devices

100.10.10.101

Results for: 100.10.10.101

- 8 hours ago: dm1-261sw11-Leaf22-DC2 on port Vlan110
- 8 hours ago: dm1-261sw12-Leaf21-DC2 on port Vlan110
- 8 hours ago: dm1-261sw11-Leaf22-DC2 on port Vlan310
- 8 hours ago: dm1-261sw12-Leaf21-DC2 on port Vlan310
- 4 hours ago: dm1-261sw13-Leaf12-DC2 on port Vlan110
- 4 hours ago: dm1-261sw13-Leaf12-DC2 on port Vlan310

Vlan110 Config

- Description:
- Forwarding Model:
- IPv4 Address:
- Loopback Mode:
- Uri Link Mode:
- Access VLAN: N/A
- Port Channel: N/A
- LACP Type: N/A
- LLDP Neighbors with Hostnames: N/A

Vlan110 on dm1-261sw11-Leaf22-DC2 Traffic Metrics

Vlan110 on dm1-261sw11-Leaf22-DC2 Error Metrics

Vlan110 Details

- Burned-in MAC Address:
- Transceiver Type:
- Duplex:
- Speed:
- Auto Negotiation Mode:
- MTU:

Device Addresses

- Mac Address: 00:10:10:00:01:01
- IP Address: 100.10.10.101

Summary

Now with Arista CloudEOS solution, customers can build the network at the ‘speed of cloud’ using repeatable design patterns powered by CloudEOS Router, deployed and provisioned by Terraform and CloudVision, and operate a consistent, secure multi-cloud environment with EOS and CloudVision across data center, campus and public clouds.

Appendix

VPC	Virtual Private Cloud
VNET	Virtual Network
TGW	Transit Gateway
CNPS	Cloud Network Private Segment
DPS	Dynamic Path Selection
DPDK	Data Plane Development Kit
VXLAN	Virtual Extensible LAN
VTEP	Virtual Tunnel Endpoint
EOS	Extensible Operating System
CVP	Cloud Vision Portal
CSP	Cloud Service Provider
IPSEC	IP Security
BGP	Border Gateway Protocol
EVPN	Ethernet Virtual Private Network
NRLI	Network Layer Reachability Information
VRF	Virtual Routing and Forwarding
RR	Route Reflector
ISP	Internet Service Provider

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390

Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2020 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. 09-0004-01 September 21, 2020