Arista Enterprise WAN design using AWS Cloud WAN

Introduction

AWS Cloud WAN is a recent innovation from AWS to provide a simplified and global WAN service built on top of their worldwide network. A good overview of this technology is here. We also paraphrase the highlights below.

AWS Cloud WAN is a managed wide-area networking (WAN) service that you can use to build, manage, and monitor a unified global network that connects resources running across your cloud and on-premises environments. It provides a central dashboard from which you can connect on-premises branch offices, data centers, and Amazon Virtual Private Clouds (VPCs) across the AWS global network. You can use simple network policies to centrally configure and automate network management and security tasks, and get a complete view of your global network.

Arista WAN Routing System

Arista's WAN Routing System including 5000 Series WAN Routing System and CloudEOS Router are designed for enterprise WAN use cases, combines cloud-grade Unified Cloud Networking offerings, enterprise-class WAN routing platforms, carrier/cloud-neutral Internet transit options, and the CloudVision® Pathfinder Service to simplify and improve customer wide area networks.

Based on Arista's EOS® software and CloudVision's cloud-based management and orchestration services, the Arista WAN Routing System delivers advanced traffic engineering, path computation, self-healing, encrypted fabric management, and predictive diagnostics and analytics from the cloud to the edge that simplify the WAN and lower WAN operating costs. Enterprise-class routing features are available in a broad series of physical, virtual, and cloud platforms – all using identical EOS software.

Arista CloudEOS is available in the AWS marketplace. Instances of it can be created and reside inside AWS.

This document illustrates how one can design and deploy an Arista enterprise WAN which includes the remote offices, central offices, AWS WAN edge (CloudEOS) deployments and their inter-region connectivity leveraging AWS Cloud WAN technology. We also discuss the attachments between the CloudEOS and AWS Cloud WAN. There are two ways to have this Connect attachment - with a GRE tunnel encapsulation and "no encapsulation" or tunnel-less. Both the Connect protocols are described in this document.

AWS Cloud WAN Terminology

AWS Virtual Private Cloud (VPC) Amazon Virtual Private Cloud (VPC) is a service that lets you launch AWS resources in a logically isolated virtual network that you define. A VPC can have many different compute and network elements. The compute elements are EC2 instances whereas the network elements include transit gateway, Cloud WAN and Arista CloudEOS.

AWS Cloud WAN consists of a Global network in AWS.

The global network includes a Core network which is the portion of the Global network managed by AWS. For example the transit gateways managed by the enterprise will be part of the global network but not the core network.

A key feature of the core network as provided by AWS is that it is truly worldwide. For example, AWS has dozens of regional Points of Presence (PoPs) spanning the US, the Americas, Europe, and Asia Pacific. A typical enterprise can leverage AWS Cloud WAN and create a worldwide without configuring a network in each region and stitching them together via links such as leased lines. No need to negotiate with individual Telcos around the world.

Core Network Edges (CNE) – these are edges of the core network in each region.

Attachments – this is the connection between the enterprise SD-WAN hub mentioned above and a CNE. This kind of attachment is called a Connect attachment, whereas a connection between a Virtual Private Cloud (VPC) and a CNE is called a VPC attachment.

Segments - a core network is partitioned into one or more segments. Segments can be thought of as distinct routing domains, Virtual Routing and Forwarding (VRF) instances or tenants. Typically, you allow communication only within a segment although the Cloud WAN provides a few ways to enable or disable inter segment communication.

AWS Cloud WAN with CloudEOS

Below we show the core network of the Global Cloud WAN. It is the portion of the Global CLoud WAN managed by AWS.

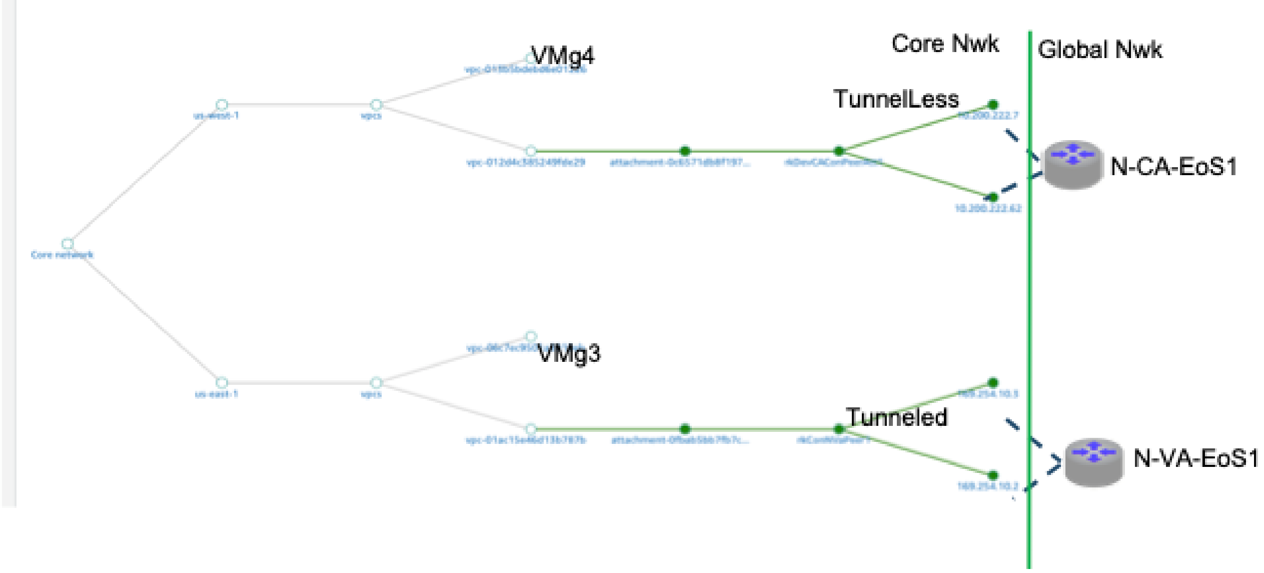

Figure 1: AWS Core Network

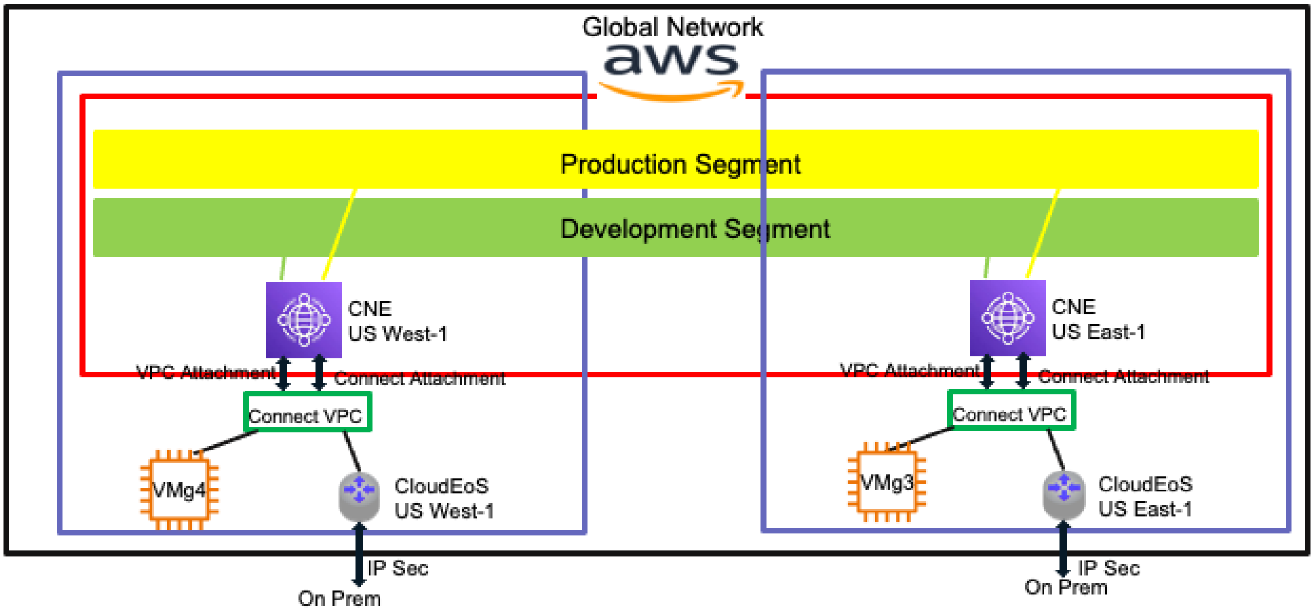

Figure 2: AWS Global Cloud WAN showing the core network and the other elements such as guest VMs and CloudEOS

In this document, N-VA is synonymous with US-East-1 and N-CA is synonymous with US-West-1.

In the above picture, one can see the AWS Global Cloud WAN, which spans regions such as US-West-1 and US-East-1. The outer black rectangle represents the AWS global network. The global network includes the network managed by AWS (i.e., the core network) and also network elements within the global network not managed by AWS such as Linux VMs (g3 and g4 above) Arista CloudEOS instances (N-VA-EOS1 and N-CA-EOS1 above). These assets are within the AWS cloud but managed by the user. The purple rectangle above represents a single region such as US-East-1 and US-West-1. We can also see the customer on-premise assets such as a branch router connected via an IP sec tunnel.

Tunnel-less and Tunnel Connects

Hitherto, AWS Cloud WAN only supported GRE tunnel Connects from the CNE (Cloud WAN Edge) to outside such as AWS CloudEOS instances. Recently AWS has developed a tunnel-less connect which has much better performance although a few unique situations and legacy deployments will have the GRE tunnel connect.

Tunnel-less Connect

In this document, we demonstrate the Arista CloudEOS integration with AWS global WAN in a tunnel based and tunnel-less cloud environment. In fact, we demonstrate the reachability between US-East-1 with GRE tunnel based connect and US-West-1 with the tunnel-less connect. We demonstrate traffic flowing end to end between CloudEOS AWS instances in the US-East-1 and US-West-1 and also to guest VMs in the two regions.

AWS Integration Steps

At the high level, the following are the steps to establish end to end communication between the CloudEOS instances in the US-East-1 and US-West-1. We use tunnel-less encapsulation in the US-West.

- Create an AWS Cloud WAN. This is global.
- Create a (different) VPC in each region and instantiate a CloudEOS instance at each end.
- Create a VPC attachment at each end between the Cloud WAN and the CloudEOS instance. The GRE tunnel attachment can span multiple segments, but the tunnel-less attachment is per segment. Thus, when using the tunnel-less attachment, you need one CloudEOS instance and one VPC per segment.

- Create a Connect attachment between a segment in the Cloud WAN and an interface in the CloudEOS instance. A segment is a VRF in CloudEOS which could even be the default VRF. The Connect attachment establishes BGP peering. The Connect attachment creation has two steps. The first is to create the attachment itself and the second is to create the BGP (Connect) peer. Connect attachment and Connect peer should be created at both US-East-1 and US-West-1.
- Create a route pointing to the CNE in the Connect VPC route-table.

The Connect attachment and Connect peer steps are different for GRE tunnel and tunnel-less cases. The figures below show these differences.

Tunnel-less Connect

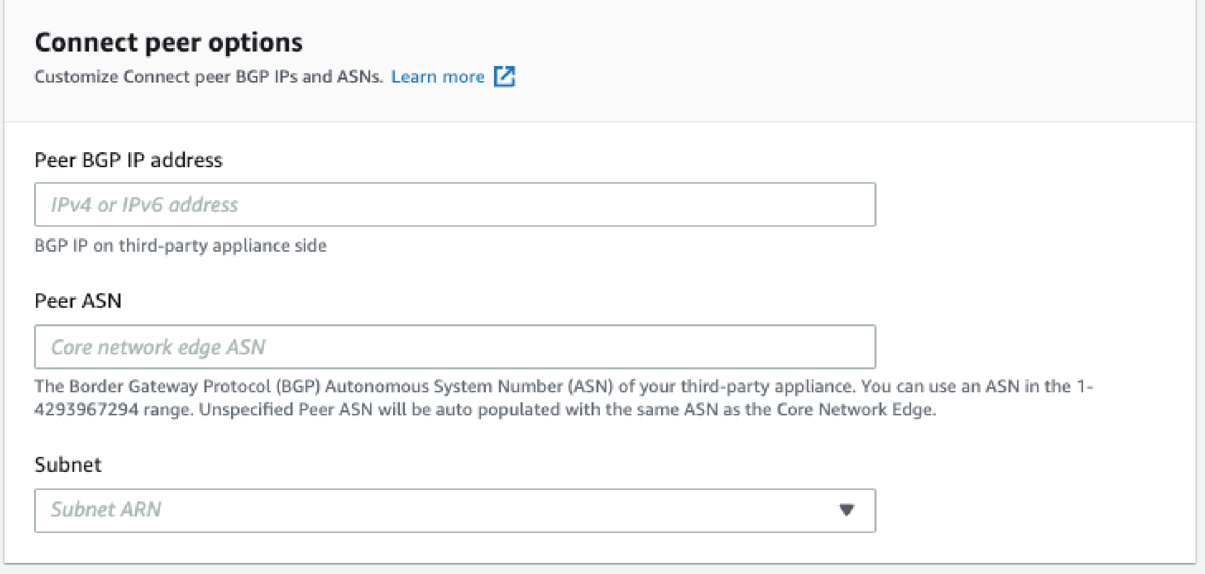

Connect peer options

Customize Connect peer BGP IPs and ASNs. Learn more

Peer BGP IP address

IPv4 or IPv6 address

BGP IP on third-party appliance side

Peer ASN

Core network edge ASN

The Border Gateway Protocol (BGP) Autonomous System Number (ASN) of your third-party appliance. You can use an ASN in the 1-4293967294 range. Unspecified Peer ASN will be auto populated with the same ASN as the Core Network Edge.

Subnet

Subnet ARN ▼

Figure 3: Tunnel-less Connect Peer

We can see above that the Connect peer for tunnel-less includes the peer BGP IP, peer ASN and subnet ARN.

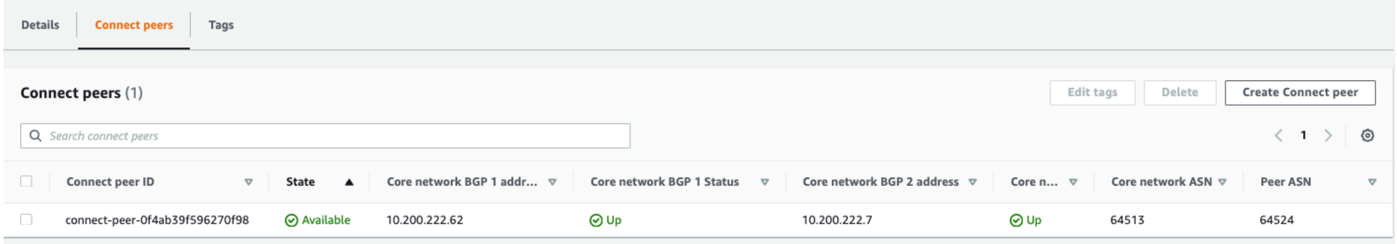

Connect peers (1)

□	Connect peer ID	State	Core network BGP 1 addr...	Core network BGP 1 Status	Core network BGP 2 address	Core n...	Core network ASN	Peer ASN
□	connect-peer-0f4ab39f596270f98	✓ Available	10.200.222.62	✓ Up	10.200.222.7	✓ Up	64513	64524

Figure 4: Tunnel-less connect in Cloud WAN

As mentioned in the AWS guidelines for tunnel-less connect, we also create a route pointing to the CNE in the tunnel-less Connect VPC route table.

Below, we show the peering interface 10.200.221.26 on CloudEOS US-West-1.

```
ncaeos1#show ip interface 10.200.221.26 brief
                                                                              Address
Interface       IP Address             Status       Protocol           MTU    Owner
--------------- ---------------------- ------------ -------------- ---------- -------
Ethernet1       10.200.221.26/28       up           up                1500
ncaeos1

ncaeos1#show runn section router bgp
router bgp 64524
   router-id 10.200.221.26
   neighbor 10.200.222.7 remote-as 64513
   neighbor 10.200.222.7 ebgp-multihop 2
   neighbor 10.200.222.62 remote-as 64513
   neighbor 10.200.222.62 ebgp-multihop 2
   redistribute connected

ncaeos1#show ip  bgp summary
BGP summary information for VRF default
Router identifier 10.200.221.26, local AS number 64524
Neighbor Status Codes: m - Under maintenance
  Neighbor      V AS           MsgRcvd   MsgSent  InQ OutQ  Up/Down State   PfxRcd PfxAcc
  10.200.222.7  4 64513          71656     84133    0    0    4d03h Estab   6      6
  10.200.222.62 4 64513          71656     84198    0    0    4d03h Estab   6      6
ncaeos1

ncaeos1#show runn int Ethernet 1
interface Ethernet1
   no switchport
   ip address 10.200.221.26/28
ncaeos1

ncaeos1#show ip route
VRF: default
Codes: C - connected, S - static, K - kernel,
       B I - iBGP, B E - eBGP, R - RIP, I L1 - IS-IS level 1,
       I L2 - IS-IS level 2, O3 - OSPFv3, A B - BGP Aggregate,
Gateway of last resort:
 S        0.0.0.0/0 [1/0] via 10.200.221.17, Ethernet1
 C        10.20.20.0/26 is directly connected, Loopback0
 B E      10.100.1.0/24 [200/100] via 10.200.221.21, Ethernet1
 B E      10.200.201.0/24 [200/100] via 10.200.221.21, Ethernet1
 B E      10.200.202.0/24 [200/100] via 10.200.221.21, Ethernet1
 B E      10.200.217.16/28 [200/100] via 10.200.221.21, Ethernet1
 C        10.200.221.16/28 is directly connected, Ethernet1
 C        10.200.221.144/28 is directly connected, Ethernet2
 B E      10.200.221.0/24 [200/100] via 10.200.221.21, Ethernet1
 B E      169.254.10.0/29 [200/100] via 10.200.221.21, Ethernet1
ncaeos1
```

Below, we show the peering interface 10.200.221.26 on CloudEOS US-West-1.

```
ncaeos1#ping 10.200.217.27
PING 10.200.217.27 (10.200.217.27) 72(100) bytes of data.
80 bytes from 10.200.217.27: icmp_seq=1 ttl=62 time=66.6 ms
```

Below shows a ping to US-East-1 tunnel end point.

```
ncaeos1#ping 169.254.10.1
PING 169.254.10.1 (169.254.10.1) 72(100) bytes of data.
80 bytes from 169.254.10.1: icmp_seq=1 ttl=62 time=64.1 ms
```

Below is a ping to US-East-1 VM. The complete details of the VM are omitted here.

```
ncaeos1#ping 10.200.202.29
PING 10.200.202.29 (10.200.202.29) 72(100) bytes of data.
80 bytes from 10.200.202.29: icmp_seq=1 ttl=125 time=63.4 ms
```

Tunnel Connect

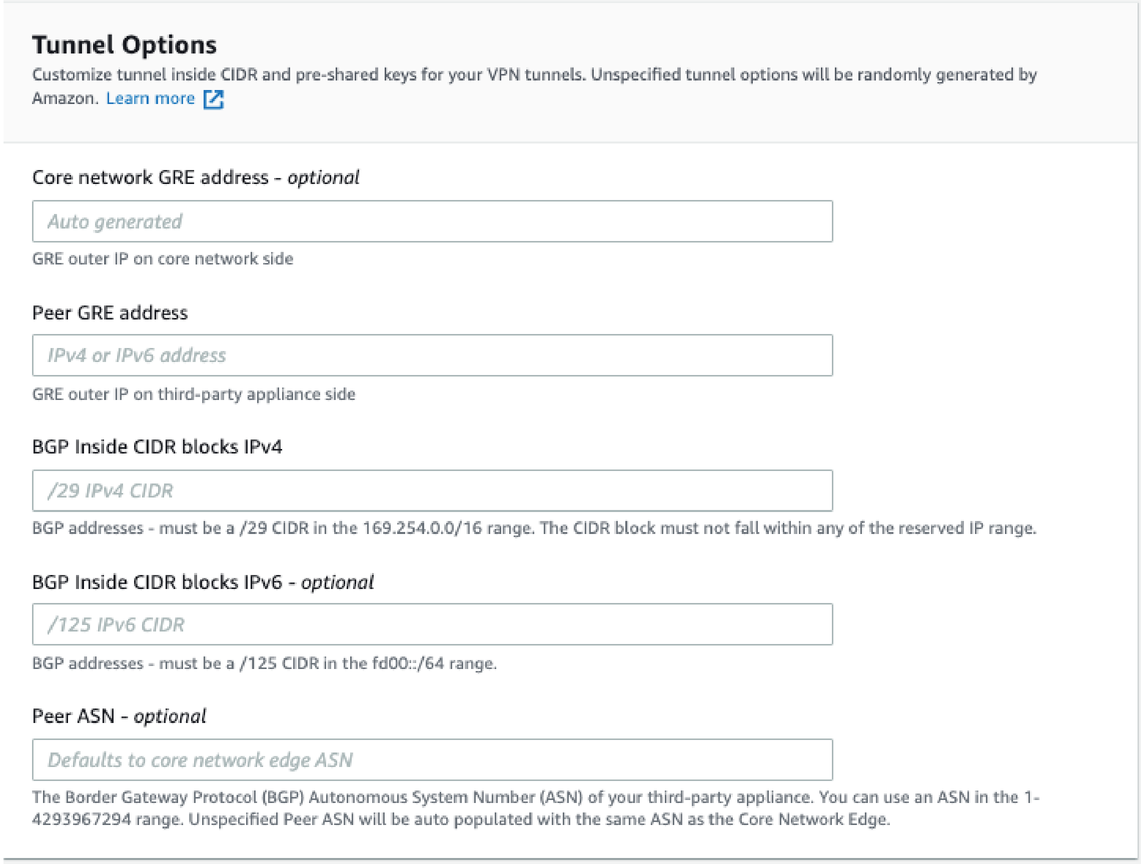

Tunnel Options

Customize tunnel inside CIDR and pre-shared keys for your VPN tunnels. Unspecified tunnel options will be randomly generated by Amazon. Learn more

Core network GRE address - *optional*

Auto generated

GRE outer IP on core network side

Peer GRE address

IPv4 or IPv6 address

GRE outer IP on third-party appliance side

BGP Inside CIDR blocks IPv4

/29 IPv4 CIDR

BGP addresses - must be a /29 CIDR in the 169.254.0.0/16 range. The CIDR block must not fall within any of the reserved IP range.

BGP Inside CIDR blocks IPv6 - *optional*

/125 IPv6 CIDR

BGP addresses - must be a /125 CIDR in the fd00::/64 range.

Peer ASN - *optional*

Defaults to core network edge ASN

The Border Gateway Protocol (BGP) Autonomous System Number (ASN) of your third-party appliance. You can use an ASN in the 1-4293967294 range. Unspecified Peer ASN will be auto populated with the same ASN as the Core Network Edge.

Figure 5: Tunnel Connect Peer

The above figure shows the GRE tunnel Connect peer. We see that for tunnel encapsulation we need the inside CIDR/IP address as well as the outer tunnel source and destination IPs as well.

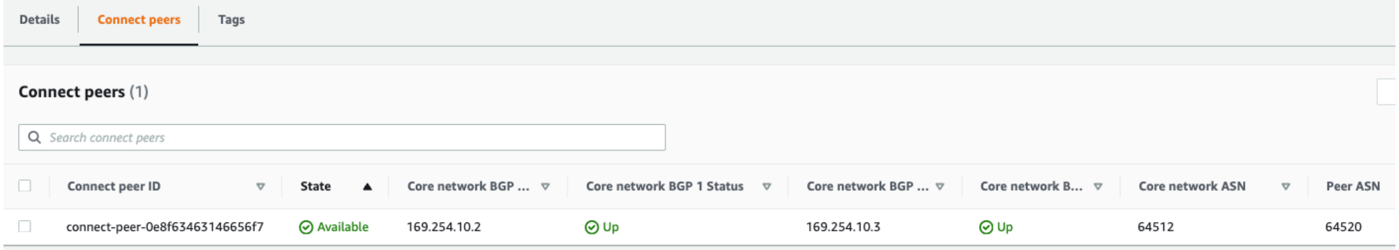

Connect peers (1)

□	Connect peer ID	State	Core network BGP ...	Core network BGP 1 Status	Core network BGP ...	Core network B...	Core network ASN	Peer ASN
□	connect-peer-0e8f63463146656f7	✓ Available	169.254.10.2	✓ Up	169.254.10.3	✓ Up	64512	64520

Figure 6: Tunnel Connect in Cloud WAN

```
nvaeos1#show ip int tunnel 0 brief
                                                                                Address
Interface        IP Address            Status       Protocol            MTU     Owner
--------------- --------------------- ------------ -------------- ---------- -------
Tunnel0          169.254.10.1/29       up           up                  1476

nvaeos1#show running-config section router bgp
router bgp 64520
   router-id 10.200.217.27
   neighbor 169.254.10.2 remote-as 64512
   neighbor 169.254.10.2 ebgp-multihop 2
   neighbor 169.254.10.3 remote-as 64512
   neighbor 169.254.10.3 ebgp-multihop 2
   network 10.100.1.0/24
   redistribute connected
nvaeos1

nvaeos1#show ip bgp summary
BGP summary information for VRF default
Router identifier 10.200.217.27, local AS number 64520
Neighbor Status Codes: m - Under maintenance
  Neighbor      V AS           MsgRcvd   MsgSent  InQ OutQ  Up/Down State   PfxRcd PfxAcc
  169.254.10.2 4 64512          563177    661175    0    0    3d01h Estab   6      6
  169.254.10.3 4 64512          541426    635946    0    0    3d02h Estab   6      6
nvaeos1

nvaeos1#show runn int tunnel 0
interface Tunnel0
   mtu 1476
   ip address 169.254.10.1/29
   tunnel source 10.200.217.27
   tunnel destination 10.200.222.23
nvaeos1
```

Below, we show all the routes learnt at US-East-1 (N. Virginia) CloudEOS router including the route to US-West-1 (N. California) CloudEOS router as well as the routes advertised from there.

```
nvaeos1#show ip route
VRF: default
Codes: C - connected, S - static, K - kernel,
       B I - iBGP, B E - eBGP, R - RIP, I L1 - IS-IS level 1,
Gateway of last resort:
 S         0.0.0.0/0 [1/0] via 10.200.217.17, Ethernet1
 B E      10.20.20.0/26 [200/100] via 169.254.10.2, Tunnel0, Static Interface GRE tunnel index 0, dst
10.200.222.23, src 10.200.217.27
 C        10.100.1.0/24 is directly connected, Loopback1
 B E      10.200.201.0/24 [200/100] via 169.254.10.2, Tunnel0, Static Interface GRE tunnel index 0, dst
10.200.222.23, src 10.200.217.27
 B E      10.200.202.0/24 [200/100] via 169.254.10.2, Tunnel0, Static Interface GRE tunnel index 0, dst
10.200.222.23, src 10.200.217.27
 C        10.200.217.16/28 is directly connected, Ethernet1
 B E      10.200.221.16/28 [200/100] via 169.254.10.2, Tunnel0, Static Interface GRE tunnel index 0, dst
10.200.222.23, src 10.200.217.27
 B E      10.200.221.144/28 [200/100] via 169.254.10.2, Tunnel0, Static Interface GRE tunnel index 0, dst
10.200.222.23, src 10.200.217.27
 B E      10.200.221.0/24 [200/100] via 169.254.10.2, Tunnel0, Static Interface GRE tunnel index 0, dst
10.200.222.23, src 10.200.217.27
 S        10.200.222.23/32 [1/0] via 10.200.217.17, Ethernet1
 C        169.254.10.0/29 is directly connected, Tunnel0, Static Interface GRE tunnel index 0, dst
10.200.222.23, src 10.200.217.27
nvaeos1
```

Thus, we see the BGP peering has been established at both ends between the CloudEOS router and the Cloud WAN.

```
nvaeos1#ping 10.200.221.26
PING 10.200.221.26 (10.200.221.26) 72(100) bytes of data.
80 bytes from 10.200.221.26: icmp_seq=1 ttl=62 time=66.0 ms
```

End to end traffic communication between CloudEOS US-East-1 and CloudEOS in US-West-1 is shown above.

Conclusion

This document summarizes the Arista Enterprise WAN deployment leveraging AWS Cloud WAN.

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, 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 © 2023 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.   October 19, 2023