



a Hewlett Packard
Enterprise company

AOS-CX IPV6 EVPN VXLAN

Contents

Introduction 3

Use Case#1 – IPv4 underlay with dual stack IPv4/IPv6 overlay (single fabric)..... 3

 IPv4 underlay 3

 Dual stack IPv4/IPv6 overlay 6

Use Case#2 – IPv6 underlay with dual stack IPv4/IPv6 overlay (single fabric)..... 8

 IPv6 underlay 8

 Dual stack IPv4/IPv6 overlay 11

References 13

 EVPN VXLAN Guide..... 13

Introduction

This document provides information on use cases supported on AOS-CX for IPv6 EVPN VXLAN.

These 2 main use cases will be covered:

- Use case#1 – IPv4 underlay with dual stack IPv4/IPv6 overlay (single fabric)
 - Recommended for existing IPv4 environments that need to support dual stack IPv4/IPv6 in the VXLAN overlay
- Use case#2 – IPv6 underlay with dual stack IPv4/IPv6 overlay (single fabric)
 - Recommended for new IPv6 environments that need to support dual stack IPv4/IPv6 in the VXLAN overlay

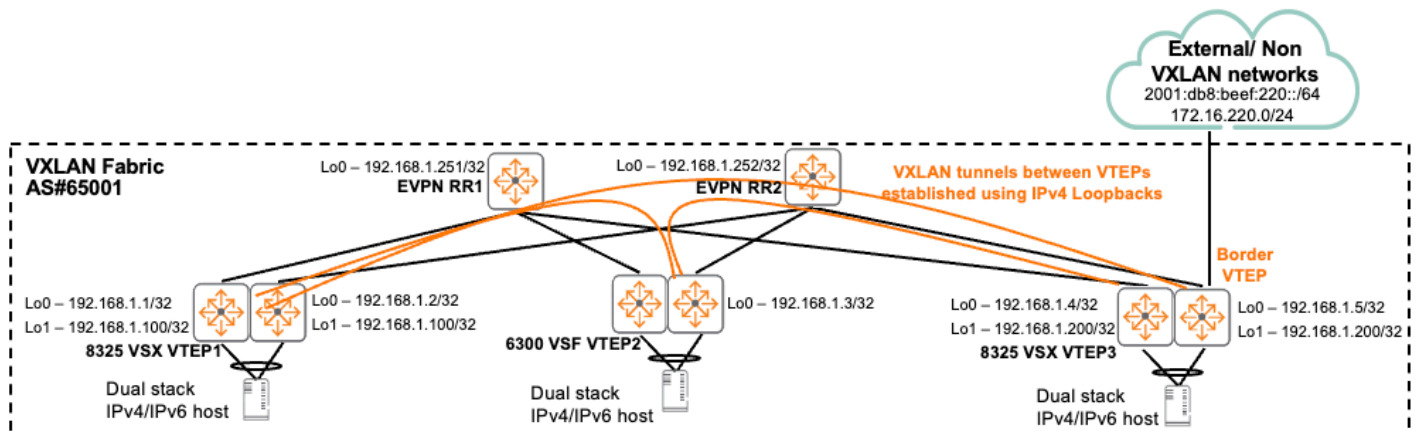
Single fabric refers to a VXLAN fabric/VXLAN domain with fully meshed VTEPs/VXLAN Tunnel End Points.

Use Case#1 – IPv4 underlay with dual stack IPv4/IPv6 overlay (single fabric)

The topology shown in Figure 1 and 2 will be used to describe this use case.

IPv4 underlay

Figure 1. Use Case#1 topology – IPv4 underlay



For the IPv4 underlay:

- VTEPs in the fabric are typically 6400/8325/8360/10000 with VSX (dual control plane) or 6300 with VSF (single control plane)
- IBGP EVPN RRs/Route Reflectors are typically 8325/8360/6400/8400/9300
- Border VTEPs are typically 8325/8360
- OSPFv2 is used to advertise and learn Loopbacks between VTEPs and RRs
- IPv4 BGP EVPN peering is established between Loopbacks on VTEPs and RRs
- VSX VTEPs utilize unique Lo0 for EVPN control plane and anycast Lo1 for data plane tunnel formation

- VSF VTEPs utilize Lo0 for EVPN control plane and data plane tunnel formation
- EVPN control plane - VTEPs do not need to peer with other VTEPs, they only need to peer with EVPN RRs
- VXLAN data plane - full mesh VXLAN tunnels between VTEPs in a fabric are formed when they interested in the same VNI/VXLAN Network Identifier

Here are config snippets for RR1.

RR2 will have a similar config with different IPs.

```
hostname RR1
!
! downlinks towards VTEPs
interface 1/1/25
no shutdown
mtu 9198
routing
ip mtu 9198
ip address 192.168.3.0/31
ip ospf 1 area 0.0.0.0
ip ospf network point-to-point
interface 1/1/26
no shutdown
mtu 9198
routing
ip mtu 9198
ip address 192.168.3.2/31
ip ospf 1 area 0.0.0.0
ip ospf network point-to-point
!
! Lo0 used for EVPN peering with VTEPs
interface loopback 0
ip address 192.168.1.251/32
ip ospf 1 area 0.0.0.0
!
!
router ospf 1
router-id 192.168.1.251
area 0.0.0.0
router bgp 65001
bgp router-id 192.168.1.251
neighbor VTEP peer-group
neighbor VTEP remote-as 65001
neighbor VTEP update-source loopback 0
neighbor 192.168.1.1 peer-group VTEP
neighbor 192.168.1.2 peer-group VTEP
neighbor 192.168.1.3 peer-group VTEP
neighbor 192.168.1.4 peer-group VTEP
neighbor 192.168.1.5 peer-group VTEP
address-family l2vpn evpn
neighbor leaf route-reflector-client
neighbor leaf send-community extended
neighbor 192.168.1.1 activate
neighbor 192.168.1.2 activate
neighbor 192.168.1.3 activate
neighbor 192.168.1.4 activate
neighbor 192.168.1.5 activate
```

Here are sample underlay configs for VTEP1.

Other VTEPs will have a similar config with different IPs.

```
hostname VTEP1

virtual-mac 02:00:00:00:01:00
! A unique virtual-mac is required on all symmetric IRB distributed L3 gateway VTEPs
! It should be the same for a VSX logical VTEP pair, different virtual-macs should be used
! for other VSX VTEP pairs
! This is the inner D-MAC/egress VTEP router MAC

vlan 10-11

! Southbound host facing ports
interface lag 1 multi-chassis
no shutdown
no routing
vlan trunk native 1
vlan trunk allowed 10-11
lacp mode active
interface 1/1/1
no shutdown
lag 1

! Northbound uplink ports
interface 1/1/23
no shutdown
mtu 9198
ip mtu 9198
ip address 192.168.3.1/31
ip ospf 1 area 0.0.0.0
ip ospf network point-to-point
interface 1/1/24
no shutdown
mtu 9198
ip mtu 9198
ip address 192.168.3.9/31
ip ospf 1 area 0.0.0.0
ip ospf network point-to-point

! Unique Lo0 on each VTEP
interface loopback 0
ip address 192.168.1.1/32
ip ospf 1 area 0.0.0.0

! Shared anycast Lo1 between a VSX logical VTEP pair
interface loopback 1
ip address 192.168.1.100/32
ip ospf 1 area 0.0.0.0

! Transit VLAN between VSX nodes for IGP continuity
interface vlan 2
ip mtu 9198
ip address 192.168.3.200/31
ip ospf 1 area 0.0.0.0
ip ospf cost 50
ip ospf network point-to-point

vsx
! The same virtual-mac and VSX system-mac are recommended for a VSX logical VTEP pair
system-mac 02:00:00:00:01:00
inter-switch-link lag 256
role primary
keepalive peer 192.168.3.3 source 192.168.3.1
```

```

router ospf 1
router-id 192.168.1.1
area 0.0.0.0


router bgp 65001
bgp router-id 192.168.1.1
neighbor RR peer-group
neighbor RR remote-as 65001
neighbor RR description RR peer-group
neighbor RR update-source loopback 0
neighbor 192.168.1.251 peer-group RR
neighbor 192.168.1.252 peer-group RR
address-family l2vpn evpn
neighbor spine-RR send-community extended
neighbor 192.168.1.251 activate
neighbor 192.168.1.252 activate
! Peer to EVPN RRs
exit-address-family
!

vrf VRF1
address-family ipv4 unicast
redistribute connected
exit-address-family
address-family ipv6 unicast
redistribute connected
! Used to advertise EVPN type 5 ip prefix routes

```

Dual stack IPv4/IPv6 overlay

Figure 2. Use Case#1 topology – dual stack IPv4/IPv6 overlay



For the dual stack IPv4/IPv6 overlay with IPv4/IPv6 unicast traffic forwarding:

- Distributed L3 Gateways are deployed on VTEPs when the same VLAN needs to be extended across multiple VTEPs, while providing optimal routing between hosts on different subnets on the same VTEP
- Hosts can utilize static IPv6 IPs, or SLAAC for dynamic IPv6 assignment
- The same link local fe80::1/64 address can be used as the IPv6 default gateway on different VLANs/SVIs as link local addresses are locally significant
- On the border VTEP, OSPF or EBGp is typically used for dynamic route exchange between the VXLAN fabric and the external/non VXLAN networks
- "Int VXLAN 1" tunnel source uses IPv4

Here are config snippets for VTEP1.

Other VTEPs will have a similar config with different IPs.

```
hostname VTEP1
!
vrf VRF1
rd 192.168.1.100:1
route-target export 65001:1 evpn
route-target import 65001:1 evpn

evpn
vlan 10
rd auto
route-target export auto
route-target import auto
redistribute host-route
vlan 11
rd auto
route-target export auto
route-target import auto
redistribute host-route
! redistribute host-route is used to advertise host routes (/32 for IPv4, /128 for IPv6)
! to remote VTEPs
! e.g. 10.1.10.100/32 and/or IPv6 address 2001:db8:1010::100/128 is connected to VTEP1
! This is required as the same subnet is spread across multiple VTEPs
! Without host routes, a remote VTEP would not be aware of the correct destination VTEP
! to send traffic

interface vlan 10
vrf attach VRF1
ip address 10.1.10.1/24
! GUA SVI IP
ipv6 address 2001:db8:1010::1/64
active-gateway ip mac 12:00:00:00:01:00
active-gateway ip 10.1.10.1
active-gateway ipv6 mac 12:00:00:00:01:00
! Link local active gateway / L3 default gateway IP
active-gateway ipv6 fe80::1

interface vlan 11
vrf attach VRF1
ip address 10.1.11.1/24
ipv6 address 2001:db8:1011::1/64
active-gateway ip mac 12:00:00:00:01:00
active-gateway ip 10.1.11.1
active-gateway ipv6 mac 12:00:00:00:01:00
active-gateway ipv6 fe80::1
! SLAAC
no ipv6 nd suppress-ra dnssl
no ipv6 nd suppress-ra rdns
! Allow RA, it is suppressed by default
ipv6 nd ra dns search-list local.net
ipv6 nd ra dns server fd25:bccf:1ffa:102::10
! Send RDNSS and DNSSL information
```

```

interface vxlan 1
! IPv4 tunnel source
source ip 192.168.1.100
no shutdown
vni 10010
vlan 10
vni 10011
vlan 11
vni 100001
routing
vrf VRF1
! L3 VNI with VRF attached

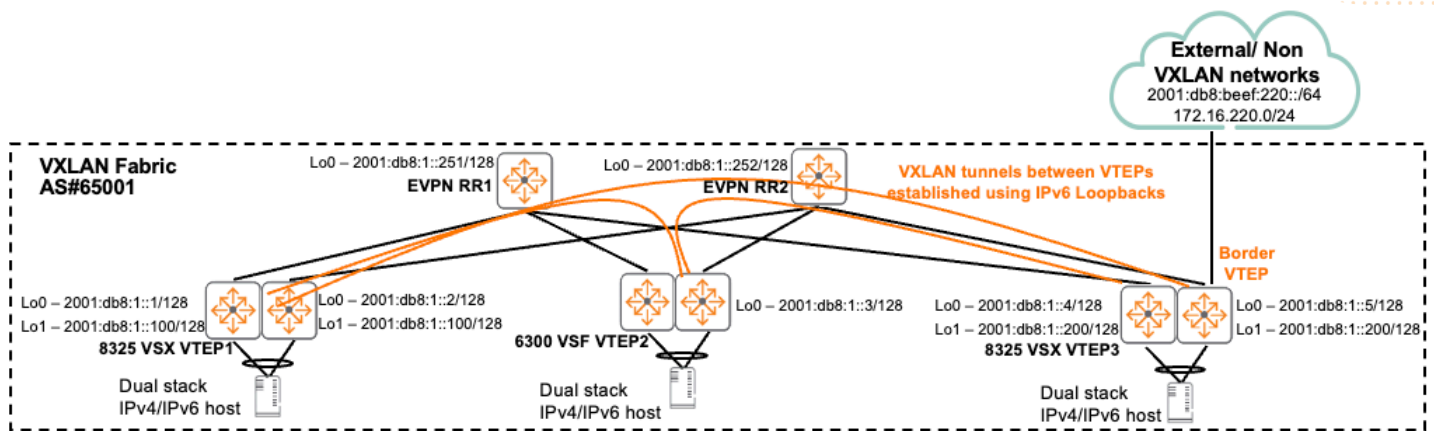
```

Use Case#2 – IPv6 underlay with dual stack IPv4/IPv6 overlay (single fabric)

The topology shown in Figure 3 and 4 will be used to describe this use case.

IPv6 underlay

Figure 3. Use Case#2 topology – IPv6 underlay



For the IPv6 underlay:

- VTEPs in the fabric are typically 6400/8360 with VSX (dual control plane) or 6300 with VSF (single control plane)
- IBGP EVPN RRs/Route Reflectors are typically 8325/8360/6400/8400/9300
- Border VTEPs are typically 8360
- OSPFv3 is used to advertise and learn Loopbacks between VTEPs and RRs
- IPv6 BGP EVPN peering is established between Loopbacks on VTEPs and RRs
- VSX VTEPs utilize unique Lo0 for EVPN control plane and anycast Lo1 for data plane tunnel formation
- VSF VTEPs utilize Lo0 for EVPN control plane and data plane tunnel formation
- EVPN control plane - VTEPs do not need to peer with other VTEPs, they only need to peer with EVPN RRs
- VXLAN data plane - full mesh VXLAN tunnels between VTEPs in a fabric are formed when they interested in the same VNI/VXLAN Network Identifier

Here are config snippets for RR1.

RR2 will have a similar config with different IPs.

```
hostname RR1
!
! downlinks towards VTEPs
interface 1/1/25
  no shutdown
  mtu 9198
  routing
  ip mtu 9198
! link local used for OSPFv3 routing
ipv6 address link-local
ipv6 ospfv3 1 area 0.0.0.0
ipv6 ospfv3 network point-to-point

interface 1/1/26
  no shutdown
  mtu 9198
  routing
  ip mtu 9198
  ipv6 address link-local
  ipv6 ospfv3 1 area 0.0.0.0
  ipv6 ospfv3 network point-to-point

!
! Lo0 used for EVPN peering with VTEPs
interface loopback 0
! IPv4 loopback required for ospfv3 and BGP router ID
ip address 192.168.1.251/32
! IPv6 GUA for EVPN VXLAN
ipv6 address 2001:db8:1::251/128
ipv6 ospfv3 1 area 0.0.0.0
!
!
router ospfv3 1
  router-id 192.168.1.1
  area 0.0.0.0

router bgp 65001
  bgp router-id 192.168.1.251
  neighbor VTEP peer-group
  neighbor VTEP remote-as 65001
  neighbor VTEP update-source loopback 0
  neighbor 2001:db8:1::1 peer-group VTEP
  neighbor 2001:db8:1::2 peer-group VTEP
  neighbor 2001:db8:1::3 peer-group VTEP
  neighbor 2001:db8:1::4 peer-group VTEP
  neighbor 2001:db8:1::5 peer-group VTEP
  address-family l2vpn evpn
  neighbor leaf route-reflector-client
  neighbor leaf send-community extended
  neighbor 2001:db8:1::1 activate
  neighbor 2001:db8:1::2 activate
  neighbor 2001:db8:1::3 activate
  neighbor 2001:db8:1::4 activate
  neighbor 2001:db8:1::5 activate
```

Here are config snippets for VTEP1.

Other VTEPs will have a similar config with different IPs.

```
hostname VTEP1
virtual-mac 02:00:00:00:01:00
! A unique virtual-mac is required on all symmetric IRB distributed L3 gateway VTEPs
! It should be the same for a VSX logical VTEP pair, different virtual-macs should be used
! for other VSX VTEP pairs
! This is the inner D-MAC/egress VTEP router MAC

vlan 10-11

! Southbound host facing ports
interface lag 1 multi-chassis
no shutdown
no routing
vlan trunk native 1
vlan trunk allowed 10-11
lacp mode active
interface 1/1/1
no shutdown
lag 1

! Northbound uplink ports
interface 1/1/23
no shutdown
mtu 9198
ip mtu 9198
ipv6 address link-local
ipv6 ospfv3 1 area 0.0.0.0
ipv6 ospfv3 network point-to-point

interface 1/1/24
no shutdown
mtu 9198
ip mtu 9198
ipv6 address link-local
ipv6 ospfv3 1 area 0.0.0.0
ipv6 ospfv3 network point-to-point

! Unique Lo0 on each VTEP
interface loopback 0
ip address 192.168.1.1/32
ipv6 address 2001:db8:1::1/128
ipv6 ospfv3 1 area 0.0.0.0

! Shared anycast Lo1 between a VSX logical VTEP pair
interface loopback 1
ipv6 address 2001:db8:1::100/128
ipv6 ospfv3 1 area 0.0.0.0

! Transit VLAN between VSX nodes for IGP continuity
interface vlan 2
ip mtu 9198
ipv6 address link-local
ipv6 ospfv3 1 area 0.0.0.0
ipv6 ospfv3 cost 50
ipv6 ospfv3 network point-to-point
```

```

vsx
! The same virtual-mac and VSX system-mac are recommended for a VSX logical VTEP pair
system-mac 02:00:00:00:01:00
inter-switch-link lag 256
role primary
keepalive peer <use OOB mgmt or dedicated KA link v6 IP> source <use OOB mgmt or dedicated
KA link v6 IP>


router ospfv3 1
router-id 192.168.1.1
area 0.0.0.0

router bgp 65001
bgp router-id 192.168.1.1
neighbor RR peer-group
neighbor RR remote-as 65001
neighbor RR description RR peer-group
neighbor RR update-source loopback 0
neighbor 2001:db8:1::251 peer-group RR
neighbor 2001:db8:1::252 peer-group RR
address-family l2vpn evpn
neighbor spine-RR send-community extended
neighbor 2001:db8:1::251 activate
neighbor 2001:db8:1::252 activate
! Peer to EVPN RRs
exit-address-family
!
vrf VRF1
address-family ipv4 unicast
redistribute connected
exit-address-family
address-family ipv6 unicast
redistribute connected
! Used to advertise EVPN type 5 ip prefix routes

```

Dual stack IPv4/IPv6 overlay

Figure 4. Use Case#2 topology – dual stack IPv4/IPv6 overlay



For the dual stack IPv4/IPv6 overlay with IPv6 unicast traffic forwarding:

- Distributed L3 Gateways are deployed on VTEPs when the same VLAN needs to be extended across multiple VTEPs, while providing optimal routing between hosts on different subnets on the same VTEP
- Hosts can utilize static IPv6 IPs or SLAAC for dynamic IPv6 assignment
- The same link local fe80::1/64 address can be used as the IPv6 default gateway on different VLANs/SVIs as link local addresses are locally significant
- On the border VTEP, OSPF or EBGp is typically used for dynamic route exchange between the VXLAN fabric and the external/non VXLAN networks
- “Int VXLAN 1 mode ipv6” tunnel source uses IPv6

Here are config snippets for VTEP1.

Other VTEPs will have a similar config with different IPs.

```
hostname VTEP1
!
vrf VRF1
rd 192.168.1.100:1
route-target export 65001:1 evpn
route-target import 65001:1 evpn

vlan 10-11

evpn
vlan 10
rd auto
route-target export auto
route-target import auto
redistribute host-route
vlan 11
rd auto
route-target export auto
route-target import auto
redistribute host-route
! redistribute host-route is used to advertise host routes (/32 for IPv4, /128 for IPv6)
! to remote VTEPs
! e.g. 10.1.10.100/32 and/or IPv6 address 2001:db8:1010::100/128 is connected to VTEP1
! This is required as the same subnet is spread across multiple VTEPs
! Without host routes, a remote VTEP would not be aware of the correct destination VTEP
! to send traffic

interface vxlan 1 mode ipv6
! IPv6 tunnel source
source ipv6 2001:db8:1::1
no shutdown
vni 10010
vlan 10
vni 10011
vlan 11
vni 100001
routing
vrf VRF1
! L3 VNI with VRF attached
```

```

interface vlan 10
vrf attach VRF1
ip address 10.1.10.1/24
! GUA SVI IP
ipv6 address 2001:db8:1010::1/64
active-gateway ip mac 12:00:00:00:01:00
active-gateway ip 10.1.10.1
active-gateway ipv6 mac 12:00:00:00:01:00
! Link local active gateway / L3 default gateway IP
active-gateway ipv6 fe80::1

interface vlan 11
vrf attach VRF1
ip address 10.1.11.1/24
ipv6 address 2001:db8:1011::1/64
active-gateway ip mac 12:00:00:00:01:00
active-gateway ip 10.1.11.1
active-gateway ipv6 mac 12:00:00:00:01:00
active-gateway ipv6 fe80::1
! SLAAC
no ipv6 nd suppress-ra dnssl
no ipv6 nd suppress-ra rdns
! Allow RA, it is suppressed by default
ipv6 nd ra dns search-list local.net
ipv6 nd ra dns server fd25:bccf:1ffa:102::10
! Send RDNSS and DNSSL information

```

References

EVPN VXLAN Guide

For additional info and verification commands, refer to

<https://www.arubanetworks.com/techdocs/AOS-CX/10.11/PDF/vxlan.pdf>

