

AOS-CX 10.6 Update  
November, 2020



# IPv6 BGP peering for EVPN AF

Aruba Switching TME



# IPv6 BGP peering for EVPN Address-Family

## Agenda



1

Overview

2

Use Cases

3

Details / Caveats

4

Configuration

5

Best Practices

6

Troubleshooting

# Overview

aruba

a Hewlett Packard  
Enterprise company

# Overview

## IPv6 BGP peering for EVPN Address-Family

- EVPN address-family is hosted in the default VRF.
- Before 10.6, BGP peering for EVPN AF could be established only between IPv4 IP addresses.
- Since 10.6, IPv6 address can be used for BGP peers to exchange EVPN routes.
- This is valid for both iBGP and eBGP peering.

# Use Cases

# Use cases

- **Phase1 of IPV6 underlay** support for EVPN-based-VXLAN fabric
- **There is not yet any use-case for this feature in 10.6.**  
Reason: only IPv4 address is supported for VXLAN tunnel address.

```
8325-1(config)# int vxlan 1
8325-1(config-vxlan-if)# source
ip IP information
8325-1(config-vxlan-if)# source ip
A.B.C.D Set the tunnel source IP
```

- When VXLAN IPv6 source address is available, then the use-case of a full IPv6 underlay for EVPN-VXLAN fabric will require such feature that is released as part of 10.6.
- Why is it in the scope of 10.6 then?  
This was a very small development effort to add this feature in 10.6, providing the necessary control-plane pre-requisite for IPv6 fabric underlay.

# Details

aruba

a Hewlett Packard  
Enterprise company

# Platform Support

10.6

Routing Protocols	6200	6300	6400	8360	8320	8325	8400
IPv6 MP-BGP peering EVPN AF	N/A	Yes	Yes	Yes	Yes	Yes	Yes
IPv6 VXLAN interface	N/A	No	No	No	No	No	No



# IPv6 BGP peering for EVPN AF - Details

- Dual-stack underlay BGP peering is not supported as not tested. IPv6 peer state might become “Connect” (verification pending):

```
8325-1# show bgp l2vpn evpn summary
VRF : default
BGP Summary
-----
Local AS           : 65001      BGP Router Identifier : 192.168.1.3
Peers              : 3          Log Neighbor Changes  : Yes
Cfg. Hold Time     : 180       Cfg. Keep Alive       : 60
Confederation Id   : 0

Neighbor      Remote-AS  MsgRcvd  MsgSent  Up/Down Time  State      AdminStatus
192.168.1.1   65001      52       674      00h:01m:25s  Established Up
fd00:192:168:1::1
               65001      0        0        00h:00m:00s  Connect   Up
```


# Configuration

aruba

a Hewlett Packard  
Enterprise company

# IPv6 BGP peering configuration

## Leaf/Spine iBGP example: OSPFv3 underlay



### Leaf

```
interface 1/1/23
  no shutdown
  mtu 9198
  ip mtu 9198
  ipv6 address fd00:192:168:3::1/127
  ipv6 ospfv3 1 area 0.0.0.0
  no ipv6 ospfv3 passive
  ipv6 ospfv3 network point-to-point
  ipv6 ospfv3 authentication ipsec spi 256 sha1 ciphertxt <snipped>
!
interface loopback 0
  ipv6 address fd00:192:168:1::3/128
  ipv6 ospfv3 1 area 0.0.0.0
!
router ospfv3 1
  router-id 192.168.1.3
  max-metric router-lsa include-stub on-startup 300
  passive-interface default
  area 0.0.0.0
```

```
leaf# show ipv6 route fd00:192:168:1::1/128

Displaying ipv6 routes selected for forwarding

'[x/y]' denotes [distance/metric]

fd00:192:168:1::1/128, vrf default, tag 0
  via fe80::883a:3001:6493:ca40%1/1/23, [110/10], ospf
```

### Spine

```
interface 1/1/25
  no shutdown
  mtu 9198
  routing
  ip mtu 9198
  ipv6 address fd00:192:168:3::0/127
  ipv6 ospfv3 1 area 0.0.0.0
  no ipv6 ospfv3 passive
  ipv6 ospfv3 network point-to-point
  ipv6 ospfv3 authentication ipsec spi 256 sha1 ciphertxt <snipped>
!
interface loopback 0
  ipv6 address fd00:192:168:1::1/128
  ipv6 ospfv3 1 area 0.0.0.0
!
router ospfv3 1
  router-id 192.168.1.1
  max-metric router-lsa include-stub on-startup 300
  passive-interface default
  area 0.0.0.0
```

```
spine# show ipv6 route fd00:192:168:1::3/128


Displaying ipv6 routes selected for forwarding

'[x/y]' denotes [distance/metric]

fd00:192:168:1::3/128, vrf default, tag 0
  via fe80::5480:2801:5cfd:f300%1/1/25, [110/10], ospf
```

# IPv6 BGP peering configuration

## Leaf/Spine iBGP example: IPv6 BGP peering



### Leaf

```
router bgp 65001
  bgp router-id 192.168.1.3
  neighbor spine-RRv6 peer-group
  neighbor spine-RRv6 remote-as 65001
  neighbor spine-RRv6 description Spine and RR peer-group for IPv6
  neighbor spine-RRv6 password ciphertext <snipped>
  neighbor spine-RRv6 fall-over
  neighbor spine-RRv6 update-source loopback 0
  neighbor fd00:192:168:1::1 peer-group spine-RRv6
  address-family l2vpn evpn
    neighbor spine-RRv6 send-community extended
    neighbor fd00:192:168:1::1 activate
  exit-address-family
```

```
leaf# show bgp l2vpn evpn summary
VRF : default
BGP Summary
-----
Local AS           : 65001      BGP Router Identifier : 192.168.1.3
Peers              : 4          Log Neighbor Changes  : Yes
Cfg. Hold Time     : 180       Cfg. Keep Alive       : 60
Confederation Id   : 0

Neighbor           Remote-AS MsgRcvd MsgSent  Up/Down Time State
AdminStatus
fd00:192:168:1::1
                  65001      49      30      00h:11m:37s Established Up
```

### Spine

```
router bgp 65001
  bgp router-id 192.168.1.1
  neighbor leafv6 peer-group
  neighbor leafv6 remote-as 65001
  neighbor leafv6 description Leaf RR Clients
  neighbor leafv6 password ciphertext <snipped>
  neighbor leafv6 fall-over
  neighbor leafv6 update-source loopback 0
  neighbor fd00:192:168:1::3 peer-group leafv6
  neighbor fd00:192:168:1::4 peer-group leafv6
  address-family l2vpn evpn
    neighbor spine-RRv6 send-community extended
    neighbor fd00:192:168:1::3 activate
  exit-address-family
```

```
spine# show bgp l2vpn evpn summary
VRF : default
BGP Summary
-----
Local AS           : 65001      BGP Router Identifier : 192.168.1.1
Peers              : 6          Log Neighbor Changes  : Yes
Cfg. Hold Time     : 180       Cfg. Keep Alive       : 60
Confederation Id   : 0

Neighbor           Remote-AS MsgRcvd MsgSent  Up/Down Time State
AdminStatus
fd00:192:168:1::3
                  65001      32      54      00h:13m:47s Established Up
fd00:192:168:1::4
                  65001      30      62      00h:14m:13s Established Up
```



aruba

a Hewlett Packard  
Enterprise company

# Best Practices

# IPv6 BGP peering for EVPN AF - Best Practices

- iBGP pering with IPv6 GUA (or ULA).
- eBGP peering with IPv6 GUA/ULA or with LLA ?
  - Currently (10.6), BGP peering with LLA is not permitted:

```
8325-1(config-bgp)# neighbor fe80::883a:3001:6493:ca40 peer-group spine-RRv6
Invalid IPV4 or IPV6 address
```

- peering with GUA/ULA address is anyhow simpler. Next-hop IPv6 address will be accepted.
  - peering with LLA address seems simpler but is actually more complex:  
(example: it might require to specify the source interface and to apply a route-map to set a valid GUA next-hop)
- All well known recommendations for IPv4 neighboring apply as well for IPv6 case:
  - fall-over, password, description, ttl-security-hops...

aruba

a Hewlett Packard  
Enterprise company

# Troubleshooting

# Troubleshooting

## Leaf/Spine iBGP example: show bgp commands

### Leaf

```
leaf# show bgp l2vpn evpn neighbors
Codes: ^ Inherited from peer-group

VRF : default

BGP Neighbor fd00:192:168:1::1 (Internal)
Description      : Spine and RR peer-group for IPv6^
Peer-group       : spine-RRv6

Remote Router Id  : 192.168.1.1      Local Router Id  : 192.168.1.3
Remote AS         : 65001            Local AS         : 65001
Remote Port       : 179              Local Port       : 39015
State            : Established       Admin Status     : Up
Conn. Established : 1                Conn. Dropped    : 0
Passive          : No                Update-Source     : loopback0^
Cfg. Hold Time   : 180              Cfg. Keep Alive  : 60
Neg. Hold Time   : 180              Neg. Keep Alive  : 60
Up/Down Time     : 00h:40m:14s      Alt. Local-AS    : 0
Local-AS Prepend : No
BFD              : Disabled
Password         : 0xaoyEjQJp2MjcFY1IIqqEbZyuMyDkoDN0zeQbI8qY0=^
Last Err Sent    : Cease
Last SubErr Sent : Connection Rejected
Last Err Rcvd    : No Error
Last SubErr Rcvd : No Error


Graceful-Restart : Enabled           Gr. Restart Time : 120
Gr. Stalepath Time : 300             Remove Private-AS : No
TTL                : 255            Local Cluster-ID :
Weight             : 0               Fall-over        : Yes^
Confederation-Peers : No

Message statistics      Sent      Rcvd
-----
Open                   1          1
Notification           2          0
Updates                20         54
Keepalives             43         36
Route Refresh          0          0
Total                  66         91

Capability              Advertised  Received
-----
Route Refresh           Yes        Yes
Graceful Restart        Yes        Yes
Add-Path                No         No
Four Octet ASN          Yes        Yes
Address family IPv4 Unicast No         No
Address family IPv6 Unicast No         No
Address family L2VPN EVPN Yes         Yes

Address Family : L2VPN EVPN

Rt. Reflect. Client : No                Send Community : extended^
```



Missing: show bgp l2vpn evpn neighbors <ipv6\_address>

### Spine

```
spine# show bgp l2vpn evpn neighbors
Codes: ^ Inherited from peer-group

VRF : default

BGP Neighbor fd00:192:168:1::3 (Internal)
Description      : Leaf RR clients for IPv6^
Peer-group       : leafv6

Remote Router Id  : 192.168.1.3      Local Router Id  : 192.168.1.1
Remote AS         : 65001            Local AS         : 65001
Remote Port       : 39015            Local Port       : 179
State            : Established       Admin Status     : Up
Conn. Established : 1                Conn. Dropped    : 0
Passive          : No                Update-Source     : loopback0^
Cfg. Hold Time   : 180              Cfg. Keep Alive  : 60
Neg. Hold Time   : 180              Neg. Keep Alive  : 60
Up/Down Time     : 00h:44m:16s      Alt. Local-AS    : 0
Local-AS Prepend : No
BFD              : Disabled
Password         : DyAc2JCw6vm8v+6PGCK7i0bZyuMyDkoDN0zeQbI8qY0=^
Last Err Sent    : Cease
Last SubErr Sent : Administrative Shutdown
Last Err Rcvd    : Cease
Last SubErr Rcvd : Connection Rejected

Graceful-Restart : Enabled           Gr. Restart Time : 120
Gr. Stalepath Time : 150             Remove Private-AS : No
TTL                : 255            Local Cluster-ID :
Weight             : 0               Fall-over        : Yes^
Confederation-Peers : No

Message statistics      Sent      Rcvd
-----
Open                   3          1
Notification           0          2
Updates                56         20
Keepalives             40         48
Route Refresh          0          0
Total                  99         71

Capability              Advertised  Received
-----
Route Refresh           Yes        Yes
Graceful Restart        Yes        Yes
Add-Path                No         No
Four Octet ASN          Yes        Yes
Address family IPv4 Unicast No         No
Address family IPv6 Unicast No         No
Address family L2VPN EVPN Yes         Yes

Address Family : L2VPN EVPN

Rt. Reflect. Client : Yes^              Send Community : extended^
```



# Troubleshooting

## Leaf/Spine iBGP example: show bgp commands

### Leaf

```
leaf# show bgp l2vpn evpn
  extcommunity    Display L2VPN EVPN routes with extended communitiescommunity
  neighbors       Display BGP neighbor connection parameters
  paths           Display the AS Path information of the routes in BGP RIB
  route-type      Displays routes filtered by NLRI route type
  summary         Summary of BGP neighbor status
  vni             VXLAN Network Identifier
  vsx-peer        Displays VSX peer switch information
  vtep            Virtual Tunnel Endpoint
  <cr>

8325-1# show bgp l2vpn evpn
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]
EVPN Route-Type 5 prefix: [5]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]
VRF : default
Local Router-ID 192.168.1.3
```

Network	Nexthop	Metric	LocPrf	Weight	Path
Route Distinguisher: 192.168.11.3:10 (L2VNI 10010)					
*> [2]:[0]:[0]:[00:50:56:8e:61:91]:[]	192.168.11.3	0	100	0	?
* i [2]:[0]:[0]:[00:50:56:8e:61:91]:[]	192.168.11.3	0	100	0	?
*> [2]:[0]:[0]:[00:50:56:8e:61:91]:[fe80::250:56ff:fe8e:6191]	192.168.11.3	0	100	0	?
* i [2]:[0]:[0]:[00:50:56:8e:61:91]:[fe80::250:56ff:fe8e:6191]	192.168.11.3	0	100	0	?
*> [2]:[0]:[0]:[12:00:00:00:01:00]:[10.1.10.1]	192.168.11.3	0	100	0	?
*> [2]:[0]:[0]:[12:00:00:00:01:00]:[fe80:10:1:10::1]	192.168.11.3	0	100	0	?
*> [3]:[0]:[192.168.11.3]	192.168.11.3	0	100	0	?
Route Distinguisher: 192.168.11.5:10 (L2VNI 10010)					
*> i [2]:[0]:[0]:[00:50:56:8e:61:91]:[10.1.10.11]	192.168.11.5	0	100	0	?
*> i [2]:[0]:[0]:[12:00:00:00:01:00]:[10.1.10.1]	192.168.11.5	0	100	0	?
*> i [3]:[0]:[192.168.11.5]	192.168.11.5	0	100	0	?

<omitted output>


### Spine

```
spine# show bgp l2vpn evpn
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]
EVPN Route-Type 5 prefix: [5]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]
VRF : default
Local Router-ID 192.168.1.1
```

Network	Nexthop	Metric	LocPrf	Weight	Path
Route Distinguisher: 192.168.11.3:10 (L2VNI 10010)					
* [2]:[0]:[0]:[00:50:56:8e:61:91]:[]	192.168.11.3	0	100	0	?
*> i [2]:[0]:[0]:[00:50:56:8e:61:91]:[]	192.168.11.3	0	100	0	?
* i [2]:[0]:[0]:[00:50:56:8e:61:91]:[]	192.168.11.3	0	100	0	?
* i [2]:[0]:[0]:[00:50:56:8e:61:91]:[fe80::250:56ff:fe8e:6191]	192.168.11.3	0	100	0	?
*> i [2]:[0]:[0]:[00:50:56:8e:61:91]:[fe80::250:56ff:fe8e:6191]	192.168.11.3	0	100	0	?
* i [2]:[0]:[0]:[00:50:56:8e:61:91]:[fe80::250:56ff:fe8e:6191]	192.168.11.3	0	100	0	?
* i [2]:[0]:[0]:[12:00:00:00:01:00]:[10.1.10.1]	192.168.11.3	0	100	0	?
*> i [2]:[0]:[0]:[12:00:00:00:01:00]:[10.1.10.1]	192.168.11.3	0	100	0	?
* i [2]:[0]:[0]:[12:00:00:00:01:00]:[10.1.10.1]	192.168.11.3	0	100	0	?
* i [2]:[0]:[0]:[12:00:00:00:01:00]:[fe80:10:1:10::1]	192.168.11.3	0	100	0	?
* i [2]:[0]:[0]:[12:00:00:00:01:00]:[fe80:10:1:10::1]	192.168.11.3	0	100	0	?
* i [3]:[0]:[192.168.11.3]	192.168.11.3	0	100	0	?
*> i [3]:[0]:[192.168.11.3]	192.168.11.3	0	100	0	?
* i [3]:[0]:[192.168.11.3]	192.168.11.3	0	100	0	?
Route Distinguisher: 192.168.11.5:10 (L2VNI 10010)					
*> i [2]:[0]:[0]:[00:50:56:8e:61:91]:[10.1.10.11]	192.168.11.5	0	100	0	?
* i [2]:[0]:[0]:[00:50:56:8e:61:91]:[10.1.10.11]	192.168.11.5	0	100	0	?
*> [2]:[0]:[0]:[00:50:56:8e:61:91]:[]	192.168.11.5	0	100	0	?
* [2]:[0]:[0]:[00:50:56:8e:61:91]:[]	192.168.11.5	0	100	0	?
*> i [2]:[0]:[0]:[12:00:00:00:01:00]:[10.1.10.1]	192.168.11.5	0	100	0	?
* i [2]:[0]:[0]:[12:00:00:00:01:00]:[10.1.10.1]	192.168.11.5	0	100	0	?
*> i [3]:[0]:[192.168.11.5]	192.168.11.5	0	100	0	?
* i [3]:[0]:[192.168.11.5]	192.168.11.5	0	100	0	?

<omitted output>



# Thank you

[vincent.giles@hpe.com](mailto:vincent.giles@hpe.com)