

CX 10.8 Update
Aug 2021

VNBT - VXLAN Stub VTEP

Presenters

- Daryl Wan
- Justin Noonan

aruba

a Hewlett Packard
Enterprise company



Agenda

- 1 Overview
- 2 Use Cases
- 3 Details and Caveats
- 3 Configuration
- 4 Best Practices
- 5 Troubleshooting
- 7 Demo
- 8 Additional Resources

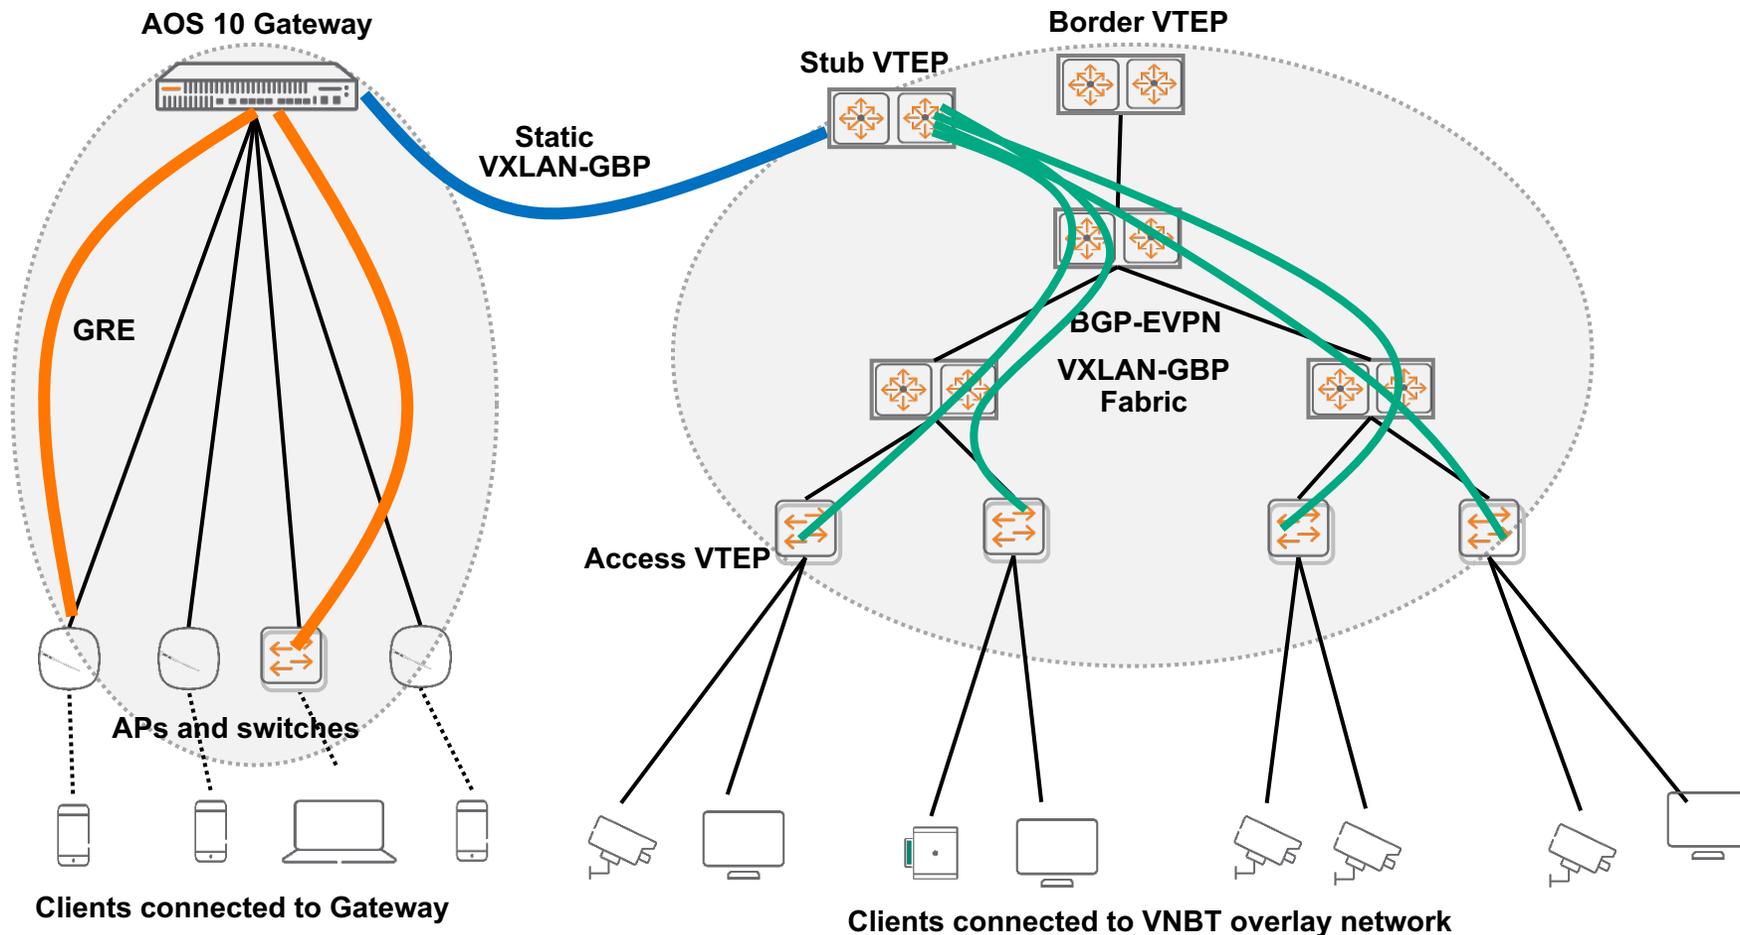




Overview

VXLAN Stub VTEP Overview

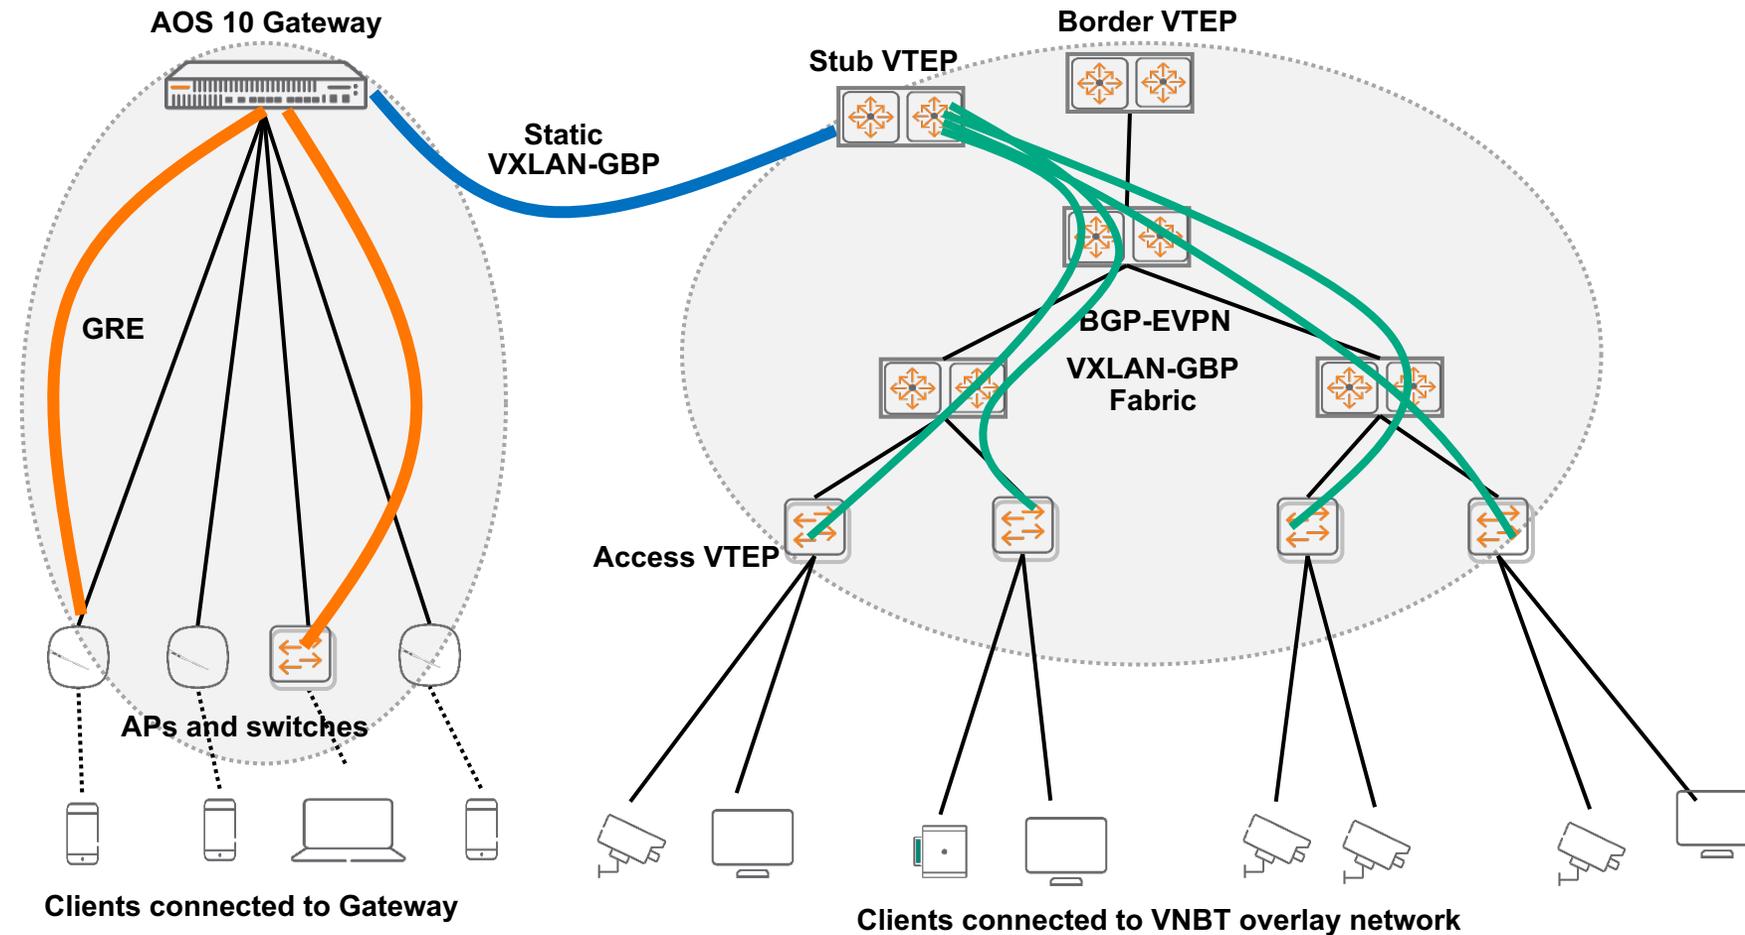
- This feature enhances the Campus VXLAN - Virtual Network Based Tunneling (VNBT) solution
- Provides VXLAN overlay network connectivity between Gateway clients and Virtual Network Based Tunneling (VNBT) clients
- Supported platforms: 6300, 6400, 8360
- Recommended platforms: 6400, 8360



The background features a solid red circle in the top-left corner and a blue area with a white dotted pattern that fills the rest of the frame.

VXLAN Stub VTEP Use Case

VXLAN Stub VTEP Use Case



Stub VTEP

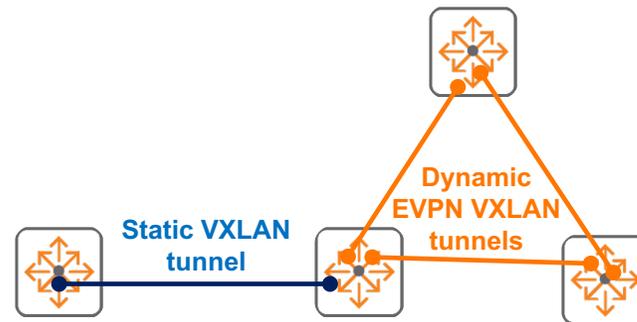
- Terminates both static VXLAN/EVPN VXLAN tunnels
- Forwards traffic between static VXLAN/EVPN VXLAN networks
- Relays VXLAN GBP between static VXLAN/EVPN VXLAN networks
- Enables role based policies between static VXLAN/EVPN VXLAN networks
- Typically functions as L3 default gateway for clients connected to gateway



VXLAN Stub VTEP Details

VXLAN Stub VTEP Details

- By default, split horizon prevents traffic (both unicast and multicast) arriving over a VXLAN tunnel from being forwarded to another VXLAN tunnel
- This is because VXLAN forwarding assumes full mesh of tunnels, allowing tunnel to tunnel forwarding will lead to loops
- In other words, VXLAN bridging across VXLAN tunnels is disabled by default



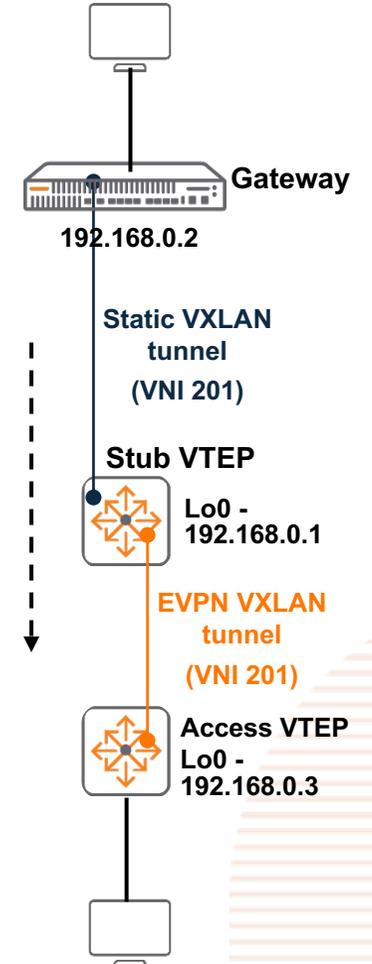
Incoming Tunnel Type	Outgoing Tunnel Type	Default Behavior of Split Horizon	Config Knob to change	Description
Dynamic	Dynamic	On	None	<ul style="list-style-type: none"> • Dynamic tunnels are fully meshed and turning it off will lead to loops
Dynamic	Static	On	inter-vxlan-bridging-mode	<ul style="list-style-type: none"> • Setting mode to static-evpn will turn off split horizon behaviour
Static	Dynamic	On	inter-vxlan-bridging-mode	<ul style="list-style-type: none"> • Setting mode to static-evpn will turn off split horizon behaviour
Static	Static	On	inter-vxlan-bridging-mode	<ul style="list-style-type: none"> • Setting mode to static-all will turn off split horizon behaviour

VXLAN GBP relay between static VXLAN/EVPN VXLAN networks

- Role (Group Policy ID) forwarded by Stub VTEP
 - From Gateway to Access VTEP

```
239... 0.499025 192.168.201.100 192.168.200.200 IPv4 10... any host intern...
▶ Frame 23903: 1066 bytes on wire (8528 bits), 1066 bytes captured (8528 bits) on interface \Device\NPF...
▶ Ethernet II, Src: ArubaaHe_b7:16:00 (64:e8:81:b7:16:00), Dst: ArubaaHe_b7:b5:00 (64:e8:81:b7:b5:00)
▶ Internet Protocol Version 4, Src: 192.168.0.2, Dst: 192.168.0.1
▶ User Datagram Protocol, Src Port: 55791, Dst Port: 4789
▼ Virtual eXtensible Local Area Network
  ▼ Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
    1... .. = GBP Extension: Defined
    .... .0.. .. = Don't Learn: False
    .... 1... .. = VXLAN Network ID (VNI): True
    .... .. 0... = Policy Applied: False
    .000 .000 0.00 .000 = Reserved(R): 0x0000
    Group Policy ID: 300
    VXLAN Network Identifier (VNI): 201
    Reserved: 0
▶ Ethernet II, Src: aa:bb:bb:00:00:02 (aa:bb:bb:00:00:02), Dst: 12:00:00:00:01:00 (12:00:00:00:01:00)
▶ Internet Protocol Version 4, Src: 192.168.201.100, Dst: 192.168.200.200
▶ Data (982 bytes)
```

192.168.201.100 aa:bb:bb:00:00:02
Role = Ultrasound (300)



192.168.200.200 aa:aa:aa:00:11:01
Role = Employee (100)

VXLAN GBP relay between static VXLAN/EVVPN VXLAN networks

- Role (Group Policy ID) forwarded by Stub VTEP
 - From Access VTEP to Gateway

```

144... 3.016055  192.168.200.200  192.168.201.100  IPv4  10... any host interna

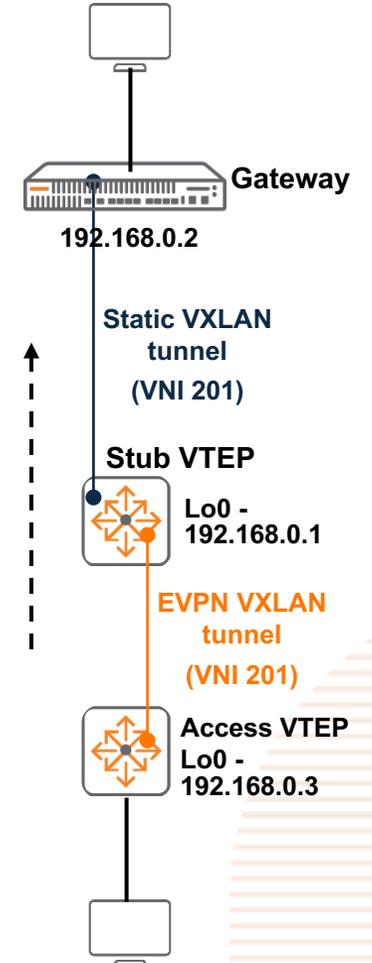
```

```

▶ Frame 144445: 1066 bytes on wire (8528 bits), 1066 bytes captured (8528 bits) on interface \Device\N
▶ Ethernet II, Src: ArubaaHe_a5:9f:40 (b8:d4:e7:a5:9f:40), Dst: ArubaaHe_b7:b5:00 (64:e8:81:b7:b5:00)
▶ Internet Protocol Version 4, Src: 192.168.0.3, Dst: 192.168.0.1
▶ User Datagram Protocol, Src Port: 35358, Dst Port: 4789
▼ Virtual eXtensible Local Area Network
  ▼ Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
    1... .. = GBP Extension: Defined
    .... ..0.. = Don't Learn: False
    .... 1... = VXLAN Network ID (VNI): True
    .... ..0... = Policy Applied: False
    .000 .000 0.00 .000 = Reserved(R): 0x0000
    Group Policy ID: 100
    VXLAN Network Identifier (VNI): 201
    Reserved: 0
▶ Ethernet II, Src: ArubaaHe_a5:9f:40 (b8:d4:e7:a5:9f:40), Dst: aa:bb:bb:00:00:02 (aa:bb:bb:00:00:02)
▶ Internet Protocol Version 4, Src: 192.168.200.200, Dst: 192.168.201.100
▶ Data (982 bytes)

```

192.168.201.100 aa:bb:bb:00:00:02
Role = Ultrasound (300)



192.168.200.200 aa:aa:aa:00:11:01
Role = Employee (100)

VXLAN Stub VTEP Caveats

- Not supported:
 - L2/L3 multicast between Static VXLAN and EVPN VXLAN on Stub VTEP is not supported
 - ARP suppression is supported only from the access ports, ARP request from VXLAN tunnels (EVPN/Static) is not supported

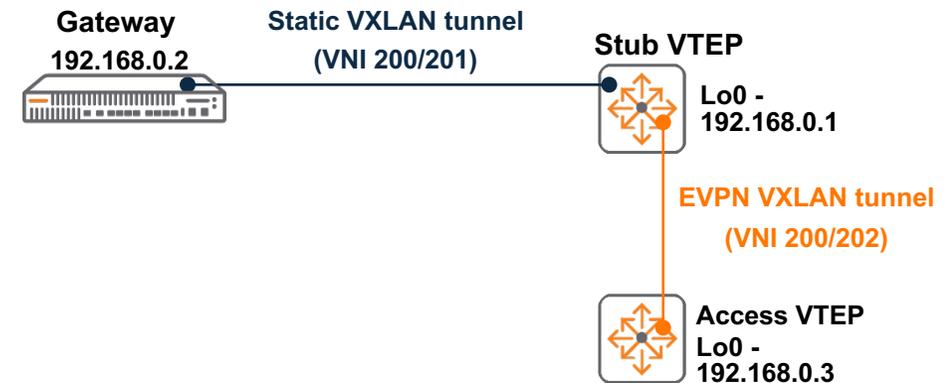


Configuration

VXLAN Stub VTEP (Distributed L3 gateway) Configuration

```
gbp enable
vlan 1,200-201
evpn
  vlan 200
    rd auto
    route-target export auto
    route-target import auto
  vlan 201
    rd auto
    route-target export auto
    route-target import auto
!
interface loopback 0
  ip address 192.168.0.1/32
  ip ospf 1 area 0.0.0.0
!
interface vxlan 1
  source ip 192.168.0.1
  inter-vxlan-bridging-mode static-evpn
  no shutdown
  vni 200
    vlan 200
    vtep-peer 192.168.0.2
  vni 201
    vlan 201
    vtep-peer 192.168.0.2
  vni 100001
    vrf VRF1
    routing
!
router ospf 1
  router-id 192.168.0.1
  area 0.0.0.0
```

Required to relay GBP tags between static/EVPN tunnels

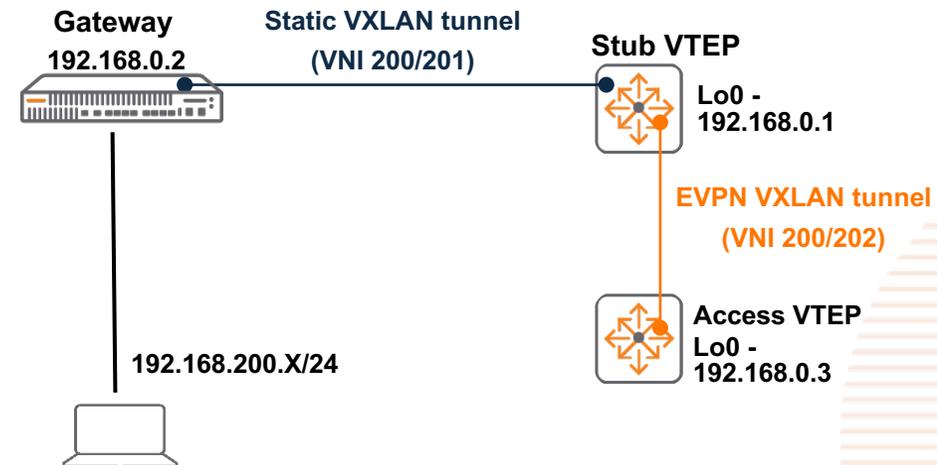


Disable split horizon between static/EVPN tunnels

VXLAN Stub VTEP (Distributed L3 gateway) Configuration

```
vrf VRF1
  rd 192.168.0.1:1
  route-target export 65100:1 evpn
  route-target import 65100:1 evpn
!
virtual-mac 02:00:00:00:01:00
!
interface vlan200
  vrf attach VRF1
  ip address 192.168.200.2/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 192.168.200.1
!
interface vlan 201
  vrf attach VRF1
  ip address 192.168.201.2/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 192.168.201.1
!
router bgp 65100
  bgp router-id 192.168.0.1
  neighbor 192.168.0.3 remote-as 65100
  neighbor 192.168.0.3 update-source loopback 0
  address-family l2vpn evpn
    neighbor 192.168.0.3 activate
    neighbor 192.168.0.3 send-community extended
```

- Typically functions as L3 default gateway for gateway clients
- Add DHCP relay if required





Best Practices

Best Practices

- Remember to enable gbp to relay GBP tags correctly between EVPN and static VXLAN tunnels
- Disable split horizon only for desired purpose
 - e.g. do not use static-all if the objective is to forward traffic between EVPN and static VXLAN tunnels
 - static-all will disable split horizon between static VXLAN tunnels too

```
interface vxlan 1
  source ip 192.168.0.1
  inter-vxlan-bridging-mode static-evpn
  no shutdown
  vni 200
    vlan 200
    vtep-peer 192.168.0.2
  vni 201
    vlan 201
    vtep-peer 192.168.0.2
```

- However, if stub VTEP is expected to forward traffic between multiple gateways, static-all should be used
- To conclude – configure according to requirements!



Troubleshooting

VXLAN Stub VTEP Troubleshooting

- Troubleshooting flow

Check underlay network reachability and VXLAN tunnels are up



Check split horizon is disabled between static and EVPN tunnels



Check MACs/IPs are learnt correctly at each VTEP

Check underlay network reachability and VXLAN tunnels are up

- In order for Stub VTEP to correctly forward traffic between static VXLAN/EVPN networks:
 - Ensure tunnel source/destination IPs are correctly advertised in the underlay network
 - e.g. “sh ip route”, test pings using loopback source/destination IPs
 - Fix underlay connectivity issues if discovered
 - If there are no underlay network issues
 - Validate both static/EVPN VXLAN tunnels

```
VNBT-Stub# sh int vxlan vteps
```

Source	Destination	Origin	Status	VNI	Routing	VLAN	VRF
192.168.0.1	192.168.0.2	static	operational	200	disabled	200	--
192.168.0.1	192.168.0.2	static	operational	201	disabled	201	--
192.168.0.1	192.168.0.3	evpn	operational	200	disabled	200	--
192.168.0.1	192.168.0.3	evpn	operational	201	disabled	201	--

- If static tunnel is down
 - Ensure correct tunnel source/destination IP is used
- If EVPN tunnel is down
 - Ensure correct EVPN configs are used

Check split horizon is disabled between static and EVPN tunnels

- If VXLAN tunnels are up, ensure split horizon is disabled between static/EVPN VXLAN tunnels

```
interface vxlan 1
 source ip 192.168.0.1
 inter-vxlan-bridging-mode static-evpn
 no shutdown
 vni 200
   vlan 200
   vtep-peer 192.168.0.2
```

```
VNBT-Stub# sh int vxlan
Interface vxlan1 is up
Admin state is up
Description:
Underlay VRF: default
Destination UDP port: 4789
VTEP source IPv4 address: 192.168.0.1
Inter vxlan bridging mode: static-evpn
```

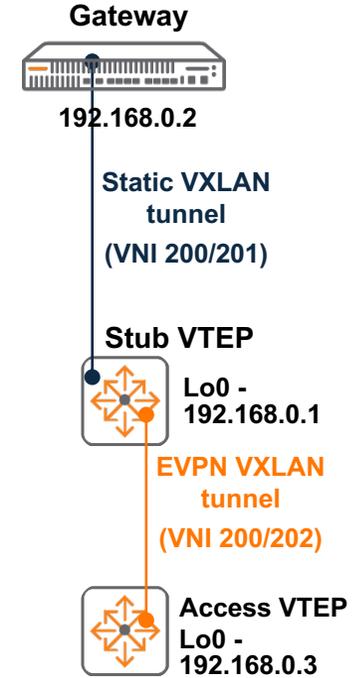
VNI	Routing	VLAN	VRF	VTEP Peers	Origin
200	disabled	200	--	192.168.0.2	static
200	disabled	200	--	192.168.0.3	evpn
201	disabled	201	--	192.168.0.2	static
201	disabled	201	--	192.168.0.3	evpn
202	disabled	202	--	192.168.0.2	static
202	disabled	202	--	192.168.0.3	evpn

Check MACs/IPs are learnt correctly at each VTEP

- On Stub VTEP, check if MACs are learnt as expected from both EVPN and Static VTEPs

```
VNBT-Stub# sh mac-ad
MAC age-time          : 300 seconds
Number of MAC addresses : 8
```

MAC Address	VLAN	Type	Port
a0:36:9f:2a:06:af	1	dynamic	1/1/1
aa:aa:aa:00:11:01	200	evpn	vxlan1 (192.168.0.3)
aa:bb:bb:00:00:02	201	dynamic	vxlan1 (192.168.0.2)
aa:cc:cc:00:00:02	201	dynamic	vxlan1 (192.168.0.2)
aa:bb:bb:00:11:02	201	evpn	vxlan1 (192.168.0.3)
aa:cc:cc:00:11:03	201	evpn	vxlan1 (192.168.0.3)



- Check if MACs are learnt as expected from both Static and EVPN VTEPs

```
Aruba-Gateway# sh mac-ad
MAC age-time          : 300 seconds
Number of MAC addresses : 6
```

MAC Address	VLAN	Type	Port
aa:aa:aa:00:11:01	200	dynamic	vxlan1 (192.168.0.1)
aa:cc:cc:00:00:02	201	port-access-security	1/1/1
aa:bb:bb:00:00:02	201	port-access-security	1/1/1
aa:bb:bb:00:11:02	201	dynamic	vxlan1 (192.168.0.1)
aa:cc:cc:00:11:03	201	dynamic	vxlan1 (192.168.0.1)

```
VNBT-Access# sh mac-ad
MAC age-time          : 300 seconds
Number of MAC addresses : 5
```

MAC Address	VLAN	Type	Port
aa:aa:aa:00:11:01	200	port-access-security	1/1/1
aa:cc:cc:00:11:03	201	port-access-security	1/1/1
aa:bb:bb:00:11:02	201	port-access-security	1/1/1
aa:bb:bb:00:00:02	201	evpn	vxlan1 (192.168.0.1)
aa:cc:cc:00:00:02	201	evpn	vxlan1 (192.168.0.1)

Check MACs/IPs are learnt correctly at each VTEP

- Check IPs are learnt as expected on both Stub (ARP table should be checked since it functions as default gateway for gateway clients) and access VTEPs

```
VNBT-Stub# show arp vrf VRF1
```

IPv4 Address	MAC	Port	Physical Port	State	VRF
10.5.8.1	5c:b9:01:1d:d5:00	vlan58	1/1/48	reachable	VRF1
10.5.8.12	00:0c:29:c5:4c:3f	vlan58	1/1/48	reachable	VRF1
10.5.8.32	2c:41:38:7f:07:06	vlan58	1/1/48	reachable	VRF1
10.5.8.3	58:20:b1:b2:9b:3f	vlan58	1/1/48	reachable	VRF1
192.168.201.12	cc:bb:bb:00:00:02	vlan201	vxlan1(192.168.0.2)	reachable	VRF1

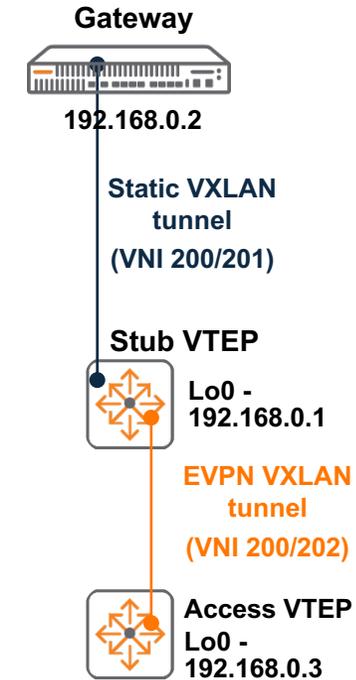
```
Total Number Of ARP Entries Listed: 5.
```

```
VNBT-Access# sh ip ro vrf VRF1
```

```
!snip
```

```
VRF: VRF1
```

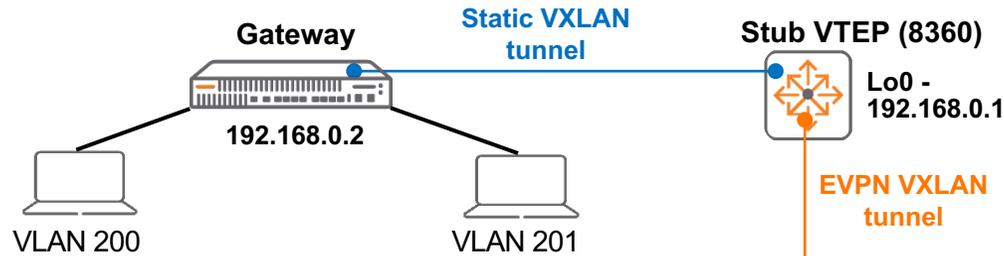
Prefix	Nexthop	Interface	VRF(egress)	Origin/Type	Distance/Metric	Age
192.168.0.0/16	192.168.0.1	-	-	B/EV	[200/0]	05d:22h:45m
192.168.10.1/32	-	loopback10	-	L	[0/0]	-
192.168.200.0/24	-	vlan200	-	C	[0/0]	-
192.168.200.1/32	-	vlan200	-	L	[0/0]	-
192.168.201.0/24	192.168.0.1	-	-	B/EV	[200/0]	05d:22h:45m
192.168.201.1/32	192.168.0.1	-	-	B/EV	[200/0]	05d:22h:45m
192.168.201.12/32	192.168.0.1	-	-	B/EV	[200/0]	00h:39m:09s
192.168.202.0/24	-	vlan202	-	C	[0/0]	-
192.168.202.1/32	-	vlan202	-	L	[0/0]	-





Demo

VXLAN Stub VTEP Demo



VLAN 200 (user) –
192.168.200.101 aa:aa:aa:00:00:01 Role = Employee (100)

VLAN 201 (device) –
192.168.201.100 aa:cc:cc:00:00:01 Role = Ultrasound (300)
192.168.201.102 cc:bb:bb:00:00:01 Role = Heart-Monitor (200)

VLAN 200 (user) –
192.168.200.201 aa:aa:aa:00:00:10 Role = Employee (100)
192.168.200.110 cc:aa:aa:00:00:02 Role = Guest (0 - default)

VLAN 202 (device) –
192.168.202.200 aa:cc:cc:00:00:10 Role = Ultrasound (300)
192.168.202.201 cc:bb:bb:00:00:10 Role = Heart-Monitor (200)

Use Case 1

- Static VXLAN tunnel from Stub to Gateway
- EVPN/VXLAN tunnel from Stub to Access

Use Case 2

- Learn MAC addresses from both static and VXLAN
- Relay GBP

The Aruba logo is displayed in a bold, orange, lowercase sans-serif font. It is positioned in the upper right corner of the slide. The background of the slide features a light blue dot grid pattern that is denser on the left and fades out towards the right. A large, curved orange shape is visible on the right side of the slide, partially overlapping the Aruba logo.

aruba

a Hewlett Packard
Enterprise company

Thank you

daryl.wan@hpe.com

justin.noonan@hpe.com