Hewlett Packard Enterprise

HPE aruba networking

AOS-CX 10.13 Update: EVPN-VXLAN with Policy Based Routing (PBR)

Corey Dow, Technical Marketing Engineer

November 11, 2023

Agenda

1	Overview
2	Use Cases
3	Details
4	Configuration
5	Best Practices
6	Troubleshooting
7	Demo

Overview

PBR (Policy Based Routing) Overview

- Why Policy Based Routing?
  - Traditional network routing is based on IPv4/IPv6 routing destinations and table lookups but does not consider the source networks
  - Certain use cases demand that both the source networks and destinations are evaluated to choose the path
    - Not necessarily the shortest path
- Examples
  - Security
    - Firewall and IPS
      - You might direct certain subnets to one or more firewalls to permit or deny access, e.g., students and faculty at a University
  - Dual ISPs
  - Load Balancers
  - Wan Optimizers
  - Service Chaining
- AOS-CX 10.01 introduced the feature for **non** EVPN-VXLAN network implementations

EVPN-VXLAN PBR (Policy Based Routing) Overview

Aruba AOS-CX 10.09 Update - EVPN VXLAN PBR

- AOS-CX Release 10.9 added PBR support for EVPN-VXLAN deployments :
- Allows L3 VTEPs to redirect traffic to desired Next Hop (NH) IP Addresses over an L2 VNI
  - Unidirectional PBR policy is applied on ingress (inbound) over a VLAN SVI interface
  - Centralized or Distributed Gateways
- Prior to 10.13 PBR could not be applied if the Next Hop (NH) was behind an L3VNI
  - For example, in Figure 1 below, a class is configured to match traffic **101.176.0.0/16** with policy **NH-Forward** that, when applied to SVI VLAN101, is supposed to forward the traffic to **NH 100.206.0.1**
  - In Figure 2, however, we can see that the PBR feature fails to apply and instead forwards the traffic to **NH 22.22.22.1** based on the routing table
  - Workarounds involved creating L2 VNIs for the VLANs where PBR implementation was needed

Figure 1

```
class ip inspect-traffic
5 match any 101.176.0.0/255.255.0.0 any count

pbr-action-list PBR-Nexthops
10 nexthop 100.206.0.1

policy NH-Forward
10 class ip inspect-traffic action pbr PBR-Nexthops

interface vlan101
   apply policy NH-Forward routed-in
   vrf attach red
   ip address 101.176.255.247/16
   active-gateway ip mac 00:10:00:00:00:00
   active-gateway ip 101.176.0.1
   ip mtu 9198
```

Figure 2

```
switch# sh pbr summary
Sequence Type Nexthop
-------------------------------------------------------
No active PBR nexthop found.
-------------------------------------------------------
! NH 100.206.0.1 is a route entry learned over a L3 VNI

switch# show ip route vrf red

Displaying ipv4 routes selected for forwarding
'[x/y]' denotes [distance/metric]

100.206.0.0/16, vrf red
via 22.22.22.1[vrf default], [200/0], bgp, encap vxlan, l3vni 10000002
100.206.0.1/32, vrf red
via 22.22.22.1[vrf default], [200/0], bgp, encap vxlan, l3vni 10000002
```

EVPN VXLAN PBR (Policy Based Routing) Overview

Platform Support

PBR and VXLAN	4100i	6000	6100	6200	6300	6400	8100	8320	8325	8360	8400	9300	10000	OVA
PBR NH over VXLAN L2VNI **(New to 10.13.0001)**	No	No	No	No	**Yes**	**Yes**	**Yes**	No	Yes	Yes	**Yes**	**Yes**	Yes	No
PBR NH over VXLAN L3VNI **(New to 10.13.0001)**	No	No	No	No	**Yes**	**Yes**	**Yes**	No	**Yes**	**Yes**	No	**Yes**	**Yes**	No

- Original PBR support for VXLAN included L2VNI on platforms 8325, 8360 and CX 10000
- Now with 10.13, PBR with L2VNI is now supported on platforms 6300 and higher, with the exception being the 8320
- With L3VNI support, the PBR policy is applied to the L3VNI context instead of the VLAN SVI interface
  - Egress policies are not supported on the 8400

Use Cases

EVPN VXLAN PBR Use Case – Distributed L3 Gateways (AOS-CX 10.09)

- Traffic to 192.168.200.0/24 from **H1** normally takes the default route on **BLF3** to reach 192.168.200.0
- PBR policy applied to **VLAN interface 20** of the ingress VTEPs
  - Redirects desired traffic through the L2VNI tunnel towards **LF1**
  - Required on all Ingress VTEPs as destination route 192.168.200.0 **points at BLF3** for all traffic
- For example
  - Redirect UDP/TCP traffic with destination port 5001 to Next Hop 31.1.1.1 connected to **LF1**
  - All other traffic towards destination **192.168.200.0** out of the tunnel continues to use **NH BLF3**

VRF	VLAN	SVI	Subnet	L2VNI
VRF1	20	20	20.1.1.0/24	10020
	21	21	21.1.1.1/30	n/a
	31	31	31.1.1.1/30	n/a

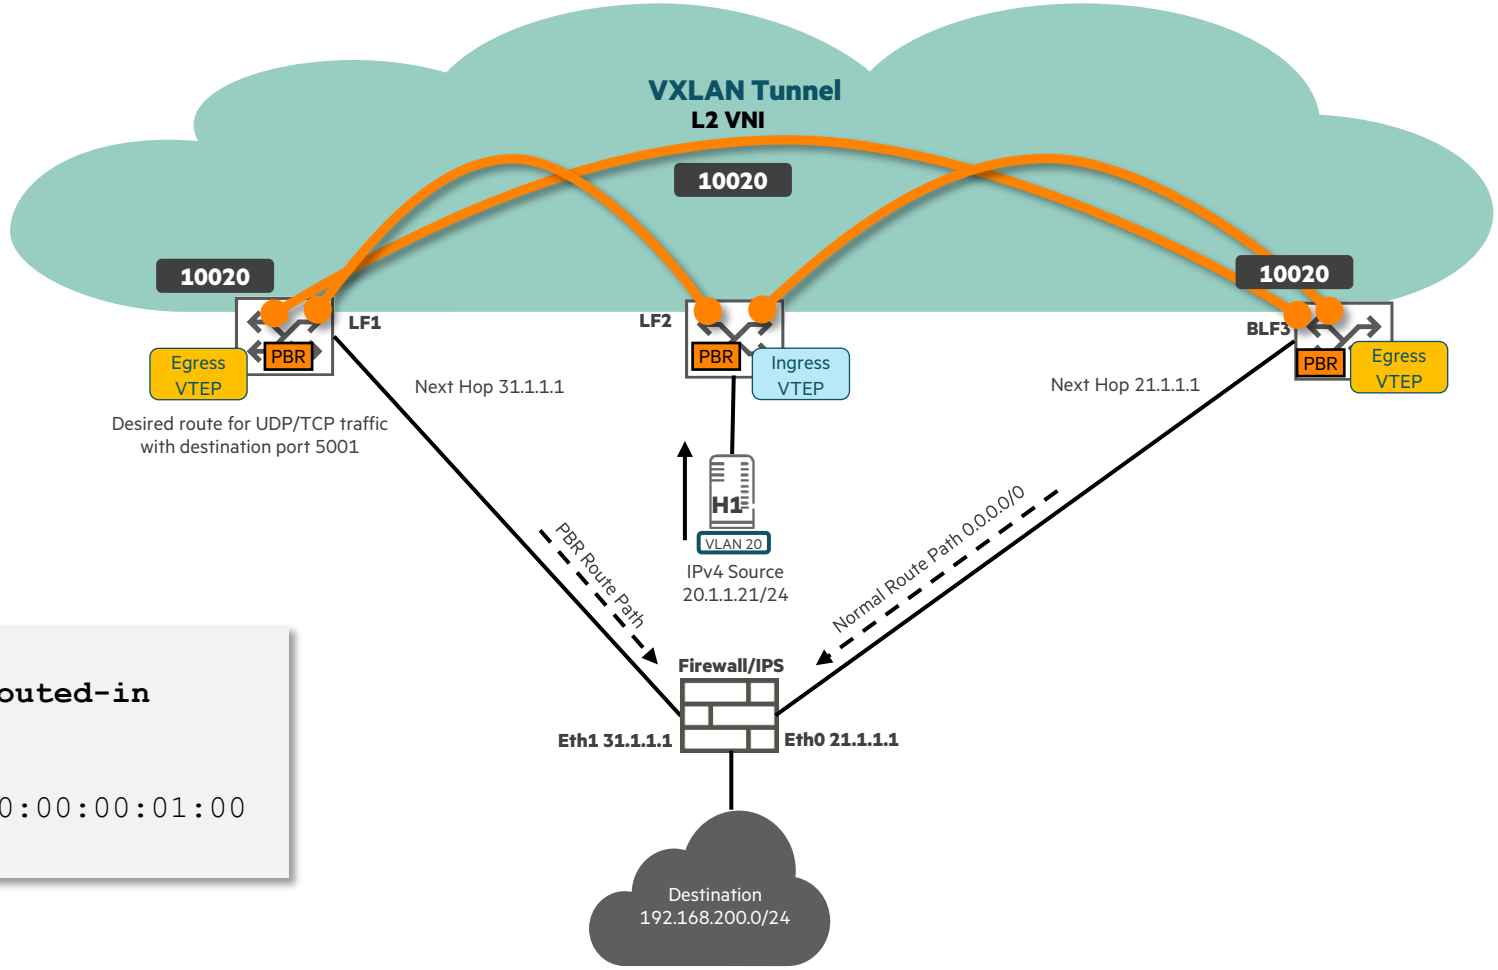

Ingress VTEP LF2

```
interface vlan 20
    apply FW_Redirect-policy routed-in
    vrf attach VRF1
    ip address 20.1.1.1/24
    active-gateway ip mac 12:00:00:00:01:00
    active-gateway ip 20.1.1.1
```

EVPN VXLAN PBR Use Case – Centralized L3 Gateways (AOS-CX 10.09)

- All traffic to 192.168.200.0/24 from **H1** and **H2** normally uses the default route on **LF1**
- Network security admins want to inspect specific traffic and redirect
- PBR policy applied to **VLAN interface 20** of ingress VTEP **LF1 only, as the gateways are centralized at LF1**
- Redirect DNS UDP traffic from both H1 and H2 to **Next Hop 21.1.1.1** on Firewall/IPS connected to **LF1**
  - To block lookups for malware sites

VRF	VLAN	SVI	Subnet	L2VNI
VRF1	20	20	20.1.1.0/24	10020
	21	21	21.1.1.1/30	n/a

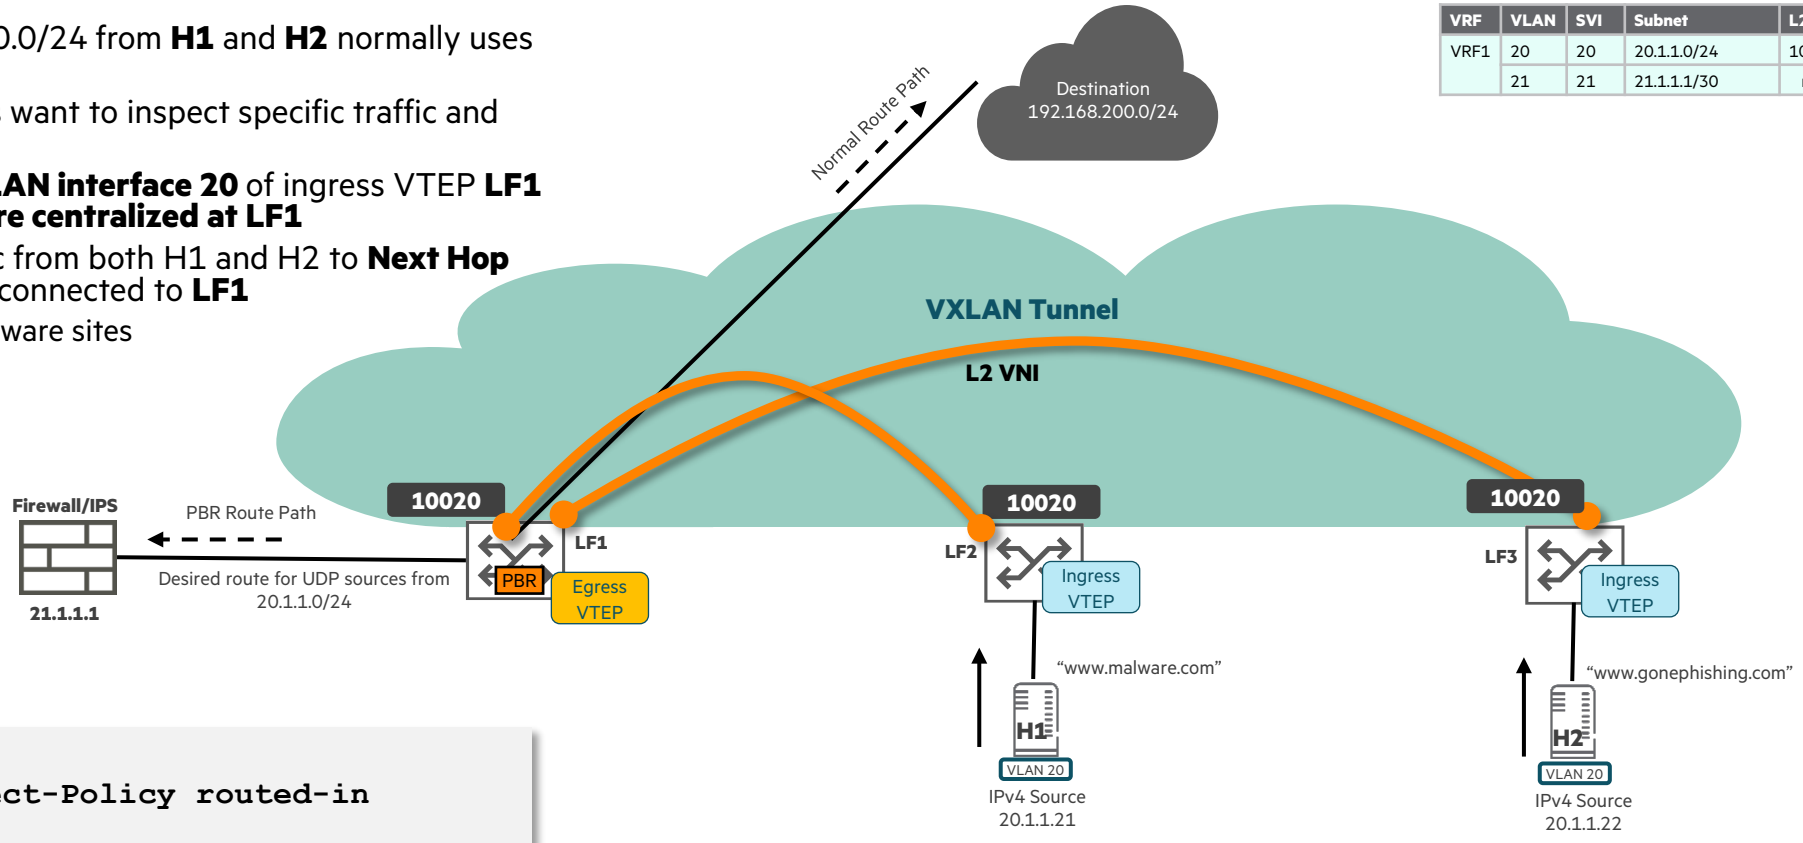

Ingress VTEP **LF1**

```
interface vlan 20
    apply FW_Redirect-Policy routed-in
    vrf attach VRF1
    ip address 20.1.1.0/24
    active-gateway ip mac 12:00:00:00:01:00
    active-gateway ip 20.1.1.1
```

EVPN VXLAN PBR Use Case – Next Hop reached over L3VNI (Distributed)

NEW

- 10.13 provides support for Next Hops reached through **an L3VNI**
  - Centralized and Distributed Anycast gateways
- Traffic to 192.168.20.20 from **H1 and H3** normally routes through the VXLAN Fabric to **LF2** as **H2** is directly connected
- Admins want to Redirect ping traffic from **H1** to **H2** to Next Hop 10.120.0.1 through **BLF3**
- PBR policy applied to **VLAN interface 10** of ingress VTEP **LF1**
  - Required to alter the normal route path to the **PBR Next Hop 10.120.0.1**
- PBR policy applied to **L3VNI 1000** of Egress VTEP **BLF3** connected directly through to the firewall off SVI 2000
  - Return traffic from the simulated firewall is routed through SVI 2001 back to **BLF3**
- All other traffic to 192.168.20.20 continues to use **LF2** as the Next Hop in the VXLAN fabric

VRF	VLAN	SVI	Subnet	L2VNI	L3VNI
VRF1	10	10	192.168.10.0/24	10	1000
	20	20	192.168.20.0/24	20	
	30	30	192.168.30.0/24	30	
	2000	2000	10.120.0.1/30	n/a	n/a
	2001	2001	10.121.0.1/30	n/a	n/a

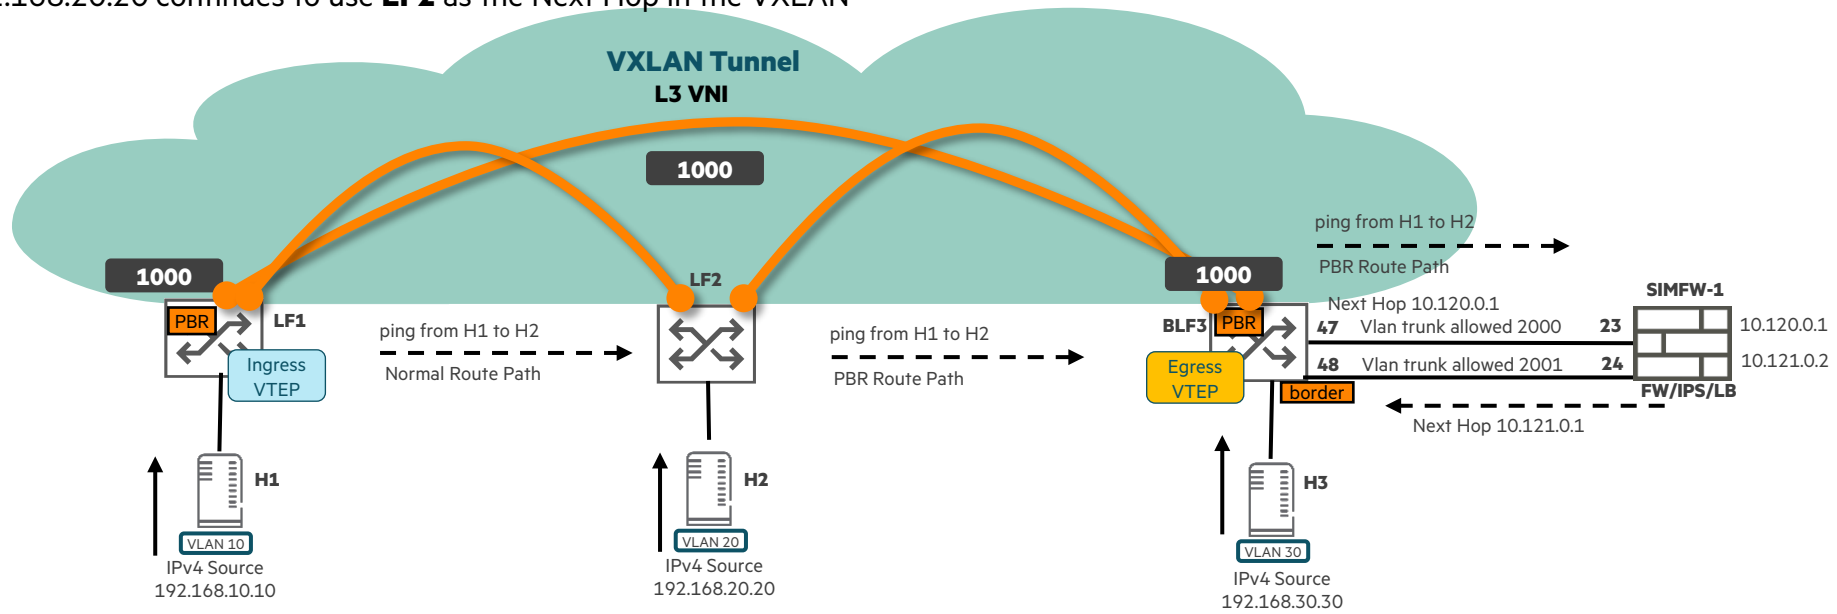

Details

EVPN VXLAN PBR Details

- Requirements
  - Applied on L3 VTEPs
    - Either centralized L3 gateway or distributed L3 gateways
  - PBR policy is applied to the ingress VTEP of an SVI or an egress VTEP with an **L3VNI** Next Hop
  - PBR Next Hop IP address must be on a subnet reachable by an Egress VTEP (SVI/ROP/L2 Switch) over an L2 or L3 VNI
- Supports
  - Next Hops that are reachable across a single fabric
    - Multi Fabric support has not yet been validated
  - Recursive Next Hops that are of type bgp evpn
    - Other route types such as BGP external are not supported
  - Class matches based on available parameters in AOS-CX:
    - UDP or TCP ports, source and destination IP ranges and DSCP values
  - Default and non-default VRFs
  - IPv4 and IPv6
  - VSX
- Specific to CX 10000
  - A security policy is required to be created in PSM for traffic to be redirected via PBR
    - PBR works with Allow ALL PSM policy
    - PBR action is performed after Elba/DPU traffic inspection
  - If traffic is not allowed via a security policy in PSM, it is dropped, and therefore cannot be redirected to PBR

Configuration

EVPN VXLAN PBR Use Case – Next Hop reached over L3VNI

NEW

- Traffic to 192.168.20.20 normally routes through the VXLAN Fabric to **LF2** as it is directly connected
- We want to redirect ping traffic **from 192.168.10.10** to destination **192.168.20.20** to Next Hop Simulated Firewall **10.120.0.1**
- All other traffic to 192.168.20.20 continues to use **LF2** as the Next Hop in the VXLAN fabric
- PBR policy applied to **VLAN interface 10** of ingress VTEP **LF1**
  - Required to alter the normal route path to **10.120.0.1 through Next Hop BLF3**
- **Next Hop 10.120.0.1 is a connected subnet**, SVI or RoP, reached through **BLF3**
- PBR policy applied to **L3VNI 1000** of Egress VTEP **BLF3** connected directly through to the firewall off **SVI 2000**
  - Return traffic from the simulated firewall is routed through **SVI 2001** back to **BLF3 tunneled to LF2**

VRF	VLAN	SVI	Subnet	L2VNI	L3VNI
VRF1	10	10	192.168.10.0/24	10	1000
	20	20	192.168.20.0/24	20	
	30	30	192.168.30.0/24	30	
	2000	2000	10.120.0.1/30	n/a	n/a
	2001	2001	10.121.0.1/30	n/a	n/a

```
VRF: red
Prefix                Nexthop                                   Interface      VRF(egress)
Origin/    Distance/    Age
Type       Metric
-----------------------------------------------------------------------------------------------
--------------
10.120.0.0/30        -                                          1/1/23          -
C          [0/0]        -
10.120.0.1/32        -                                          1/1/23          -
L          [0/0]        -
192.168.0.0/16      10.121.0.1                                 1/1/24          -
S          [1/0]        00h:24m:42s
```

Ingress VTEP

```
interface vlan 10
    vrf attach red
    ip address 192.168.10.2/24
    active-gateway ip mac aa:bb:cc:00:00:10
    active-gateway ip 192.168.10.1
    apply policy FW_Redirect routed-in
interface vxlan 1
    source ip 1.1.1.1
    no shutdown
    vni 1000
        vrf VRF1
        routing
```

Egress VTEP

```
interface vxlan 1
    source ip 3.3.3.3
    no shutdown
    vni 1000
        vrf VRF1
        routing
        apply policy FW_redirect routed-in
```

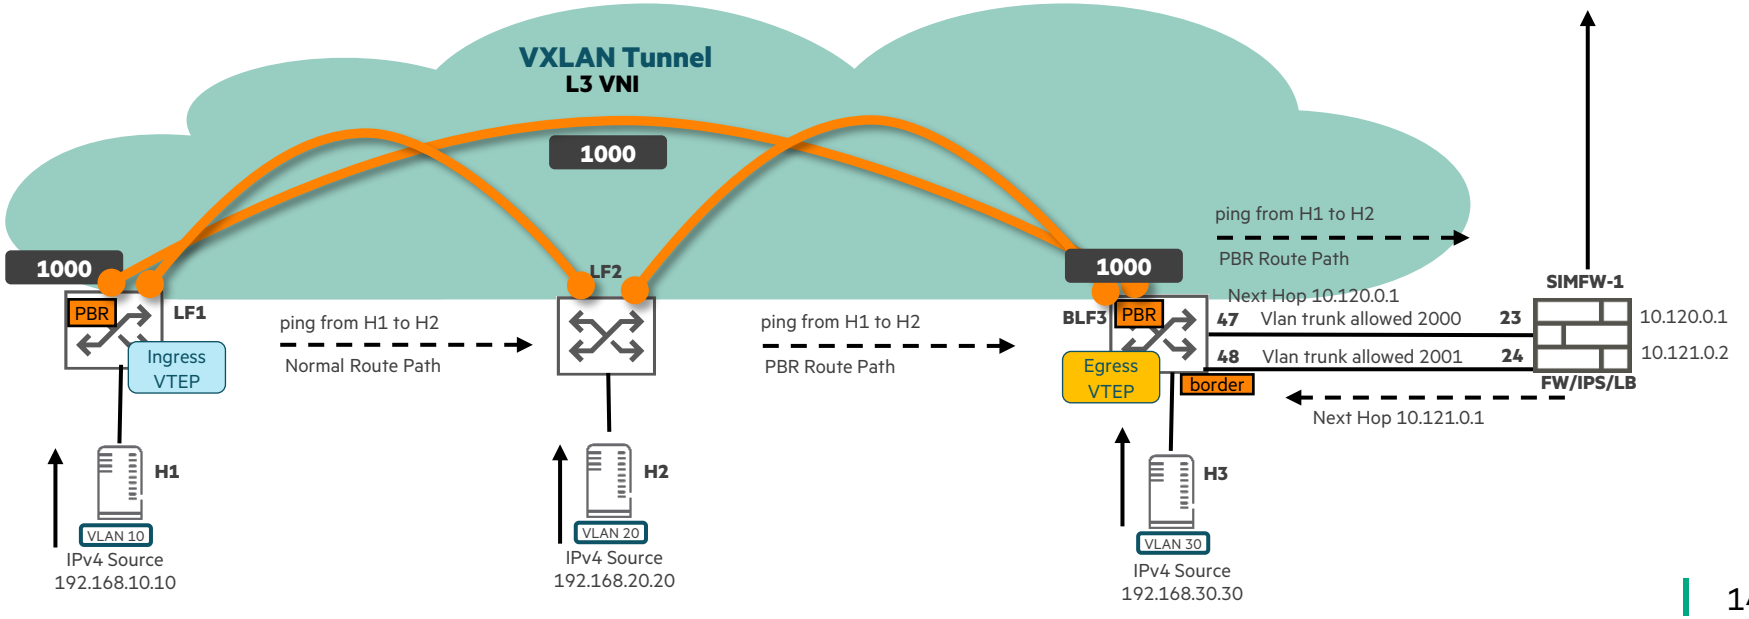

EVPN VXLAN PBR Use Case – Next Hop reached over L3VNI (IPv4)

NEW

- PBR traffic classes are created matching the desired traffic, i.e., icmp traffic **from 192.168.10.10** to **192.168.20.20**
- PBR action lists are created to identify the PBR Next Hop destination IP Addresses
- A policy of type pbr is created to pair the traffic class with the Next Hop actions
- Policies are applied to **LF1** and **BLF3** but not **LF2**

VRF	VLAN	SVI	Subnet	L2VNI	L3VNI
VRF1	10	10	192.168.10.0/24	10	1000
	20	20	192.168.20.0/24	20	
	30	30	192.168.30.0/24	30	
	2000	2000	10.120.0.1/30	n/a	n/a
	2001	2001	10.121.0.1/30	n/a	n/a

Ingress VTEP **LF1**

```
interface vlan 10
    vrf attach red
    ip address 192.168.10.2/24
    active-gateway ip mac aa:bb:cc:00:00:10
    active-gateway ip 192.168.10.1
    apply policy FW_Redirect routed-in
interface vxlan 1
    source ip 1.1.1.1
    no shutdown
    vni 1000
        vrf VRF1
        routing
```

Egress VTEP **BLF3**

```
interface vxlan 1
    source ip 3.3.3.3
    no shutdown
    vni 1000
        vrf VRF1
        routing
        apply policy FW_Redirect routed-in
```

LF2

```
interface vlan 20
    vrf attach red
    ip address 192.168.20.2/24
    active-gateway ip mac aa:bb:cc:00:00:20
    active-gateway ip 192.168.20.1
interface vxlan 1
    source ip 2.2.2.2
    no shutdown
    vni 1000
        vrf VRF1
        routing
```

PBR

```
class ip pbr-class
    10 match icmp 192.168.10.10 192.168.20.20
pbr-action-list pbr-nexthops
    10 nexthop 10.120.0.1
policy FW_Redirect
    10 class ip pbr-class action pbr pbr-nexthops
```

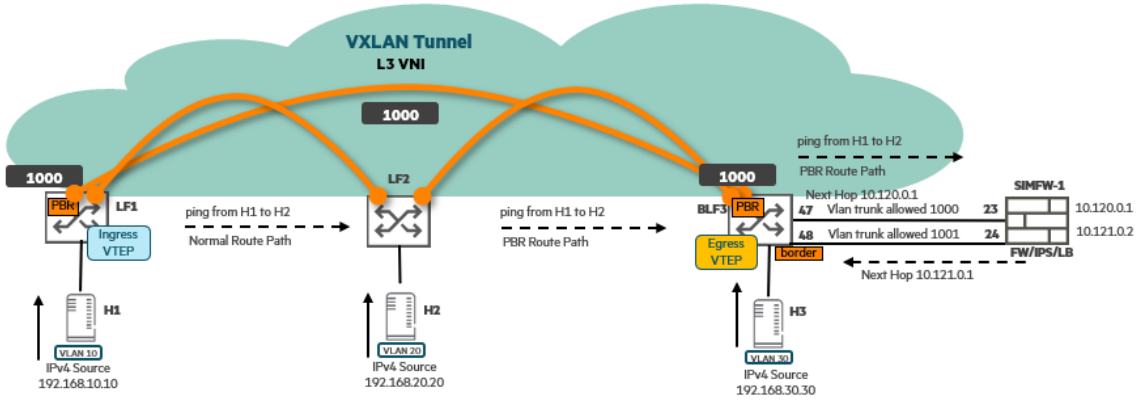

EVPN VXLAN PBR Use Case – Next Hop reached over L3VNI (IPv6)

NEW

- PBR traffic classes are created matching the desired traffic, i.e., icmp traffic from **2001:db8:192:168:10:10** to **2001:db8::192:168:20:20**
- PBR action lists are created to identify the PBR Next Hop destination IPv6 Addresses
- A policy of type pbr is created to pair the traffic class with the Next Hop actions and applied to **LF1** and **LF3**, not LF2

VRF	VLAN	SVI	Subnet	L2VNI	L3VNI
VRF1	10	10	2001:db8::192:168:10:0/64	10	1000
	20	20	2001:db8::192:168:20:0/64	20	
	30	30	2001:db8::192:168:30:0/64	30	
	2000	2000	2001:db8::10:120:0:1/127	n/a	n/a
	2001	2001	2001:db8::10:121:0:1/127	n/a	n/a

Ingress VTEP

LF1

```
interface vlan 10
    vrf attach red
    ip address 192.168.10.2/24
    active-gateway ip mac aa:bb:cc:00:00:10
    active-gateway ip 192.168.10.1
    ipv6 addresss 2001:db8::192:168:10:2/64
    active-gateway ipv6 mac aa:bb:cc:00:00:10
    active-gateway ipv6 2001:db8::192:168:10:1
    apply policy FW_Redirect routed-in

interface vxlan 1
    source ip 1.1.1.1
    no shutdown
    vni 1000
        vrf VRF1
        routing
```

Egress VTEP

BLF3

```
interface vxlan 1
    source ip 3.3.3.3
    no shutdown
    vni 1000
        vrf VRF1
        routing
        apply policy FW_Redirect routed-in
```

PBR

```
class ipv6 pbr-class
    10 match icmp 2001:db8::192:168:10:10 2001:db8::192:168:20:20
pbr-action-list pbr-nexthops
    10 nexthop 2001:db8::10:120:0:1
policy FW_Redirect
    10 class ip pbr-class action pbr pbr-nexthops
```

LF2

```
interface vlan 20
    vrf attach red
    ip address 192.168.20.2/24
    active-gateway ip mac aa:bb:cc:00:00:20
    active-gateway ip 192.168.20.1
    ipv6 address 2001:db8::192:168:20:2/64
    active-gateway ipv6 mac aa:bb:cc:00:00:20
    active-gateway ipv6 2001:db8::192:168:20:1

interface vxlan 1
    source ip 2.2.2.2
    no shutdown
    vni 1000
        vrf VRF1
        routing
```

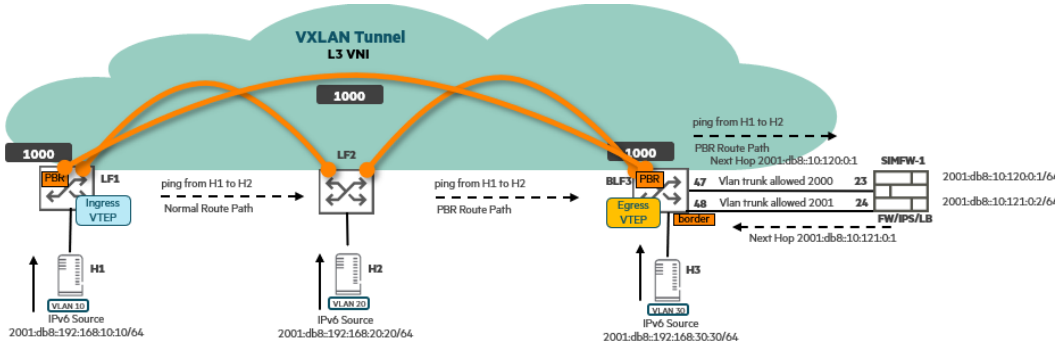

Equivalent Configurations

AOS-CX configuration comparisons

AOS-CX

Ingress Example

```
class ip ingress_class
 10 match any 200.1.1.3 30.1.1.1 class
ipv6 v6ingress_class
 10 match any 200:1:1::3/64 30:1:1::1/64

pbr-action-list ingress_pbr
 10 nexthop 100.1.1.4
 20 nexthop 201.1.1.2
 30 interface null
pbr-action-list v6ingress_pbr
 10 nexthop 100:1:1::4

policy ingress_policy
 10 class ip ingress_class action pbr
ingress_pbr
 20 class ipv6 v6ingress_class action pbr
v6ingress_pbr

interface vlan 20
    vrf attach red
    apply policy ingress_policy routed-in
```

Egress Example

```
class ip egress_class
 10 match any 200.1.1.3 30.1.1.1
class ipv6 v6egress_class
 10 match any 200:1:1::3/64 30:1:1::1/64

pbr-action-list egress_pbr
 10 nexthop 100.1.1.4
Pbr-action-list v6egress_pbr
 10 nexthop 100:1:1::4

policy l3vni_policy
 10 class ip egress_class action pbr
egress_pbr
 20 class ipv6 v6egress_class action pbr
v6egress_pbr

interface vxlan 1
source ip 192.168.0.2
vni 1000
        vrf red
        routing
        apply policy l3vni_policy routed-in
```

Cisco-NXOS

Ingress Example

```
feature pbr
ipv6 access-list bummy
  statistics per-entry
  10 permit ipv6 2001:10:1:1::20/128 any
ipaccess-list dummy
  statistics per-entry
  10 permit ip 10.1.1.20/32 any

route-map bummy permit 10
  match ipv6 address bummy
  set ipv6 next-hop 2001::DB8:800:200C:417A
 <== next-hop host behind some intermediate
VTEP

route-map dummy permit 10
  match ipaddress dummy
  set ip next-hop 10.1.1.40
 <== next-hop host behind some intermediate
VTEP

interface Vlan10
  ip policy route-map dummy
  ipv6 policy route-map bummy
```

Egress Example

```
feature pbr
ipv6 access-list bummy
  statistics per-entry
  10 permit ipv6 2001:10:1:1::20/128 any
ipaccess-list dummy
  statistics per-entry
  10 permit ip 10.1.1.20/32 any

route-map bummy permit 10
  match ipv6 address bummy
  set ipv6 next-hop 2001::DB8:800:200C:417A
 <== next-hop host behind some intermediate
VTEP

route-map dummy permit 10
match ipaddress dummy
set ip next-hop 10.1.1.40
<== next-hop host behind some intermediate
VTEP

interface Vlan2500
  ip policy route-map dummy
  ipv6 policy route-map bummy
```

- Similar configuration when compared with Cisco
  - Cisco uses route maps to set the Next Hop and uses access lists to match on entries
  - AOS-CX uses classes and policies

Equivalent Configuration

AOS-CX configuration comparisons

AOS-CX

Ingress Example

```
class ip ingress_class
 10 match any 200.1.1.3 30.1.1.1 class
ipv6 v6ingress_class
 10 match any 200:1:1::3/64 30:1:1::1/64

pbr-action-list ingress_pbr
 10 nexthop 100.1.1.4
 20 nexthop 201.1.1.2
 30 interface null
pbr-action-list v6ingress_pbr
 10 nexthop 100:1:1::4

policy ingress_policy
 10 class ip ingress_class action pbr
ingress_pbr
 20 class ipv6 v6ingress_class action pbr
v6ingress_pbr

interface vlan 20
    vrf attach red
    apply policy ingress_policy routed-in
```

Egress Example

```
class ip egress_class
 10 match any 200.1.1.3 30.1.1.1
class ipv6 v6egress_class
 10 match any 200:1:1::3/64 30:1:1::1/64

pbr-action-list egress_pbr
 10 nexthop 100.1.1.4
Pbr-action-list v6egress_pbr
 10 nexthop 100:1:1::4

policy l3vni_policy
 10 class ip egress_class action pbr
egress_pbr
 20 class ipv6 v6egress_class action pbr
v6egress_pbr

interface vxlan 1
source ip 192.168.0.1
vni 1000
        vrf red
        routing
        apply policy l3vni_policy routed-in
```

H3C/HPE comware

Egress Example

```
interface vsi-interface 3
 ip binding vpn-instance vpna
 l3-vni 1000

acl advanced 3000
 rule 0 permit ip source 10.1.1.10 0
destination 10.1.1.20 0

policy-based-route aaa permit node 0
 if-match acl 3000
 apply service-chain path-id 1
 apply next-hop vpn-instance vpna 10.1.1.11

interface vsi-interface 3
 ip policy-based-route aaa
```

- Similar configuration when compared with H3C/HPE comware
  - Comware uses ACLs and policy-based-route policies to set the Next Hop
  - AOS-CX uses classes and policies

Best Practices

Best Practices

- Add **"count"** keywords to configured pbr classes to verify the traffic entries are being hit
- In the example below, we can see that class entries 30 and 40 are getting hit with traffic from subnets 172.182.0.0/16 and 172.185.0.0/16

```
Bldg1-Acc-6300-VSF# show policy hitcounts pbr-inspect interface vlan20
Statistics for Policy pbr-inspect:

VRF VRF1
interface vlan 20 (routed-in):
     Matched Packets  Configuration
10 class ip pbr-class action pbr pbr-nexthops
                   0  10 ignore any 17.181.1.101 any count
                   0  20 match any 17.181.0.0/255.255.0.0 any count
                  12  30 match any 17.182.0.0/255.255.0.0 any count
                 100  40 match any 17.185.0.0/255.255.0.0 any count
```

- Avoid sequencing class entries sequentially, i.e., 1, 2, 3, 4 in the event you need to configure additional entries later that match before the next sequence

```
class ip pbr-class
    10 ignore any 17.181.1.101 any count
    20 match any 17.181.0.0/255.255.0.0 any count
    30 match any 17.182.0.0/255.255.0.0 any count
    40 match any 17.185.0.0/255.255.0.0 any count
```

Troubleshooting

EVPN-VXLAN PBR Troubleshooting

- Have a topology diagram ready with IP Addresses and interface details
- Recommended troubleshooting flow:

1. Check VXLAN PBR configurations are correctly applied
2. Check policy hitcounts to verify traffic classes are matching
3. Check Next Hop reachability via L2/L3 VNI
4. Verify traffic is sent to the Next Hop IP (capture)
5. Check for L3 Forwarding Loops – TTL exceeded
6. Check resources (TCAM exhaustion)

VRF	VLAN	SVI	Subnet	L2VNI	L3VNI
VRF1	10	10	192.168.10.0/24	10	1000
	20	20	192.168.20.0/24	20	
	30	30	192.168.30.0/24	30	
	2000	2000	10.120.0.1/30	n/a	n/a
	2001	2001	10.121.0.1/30	n/a	n/a

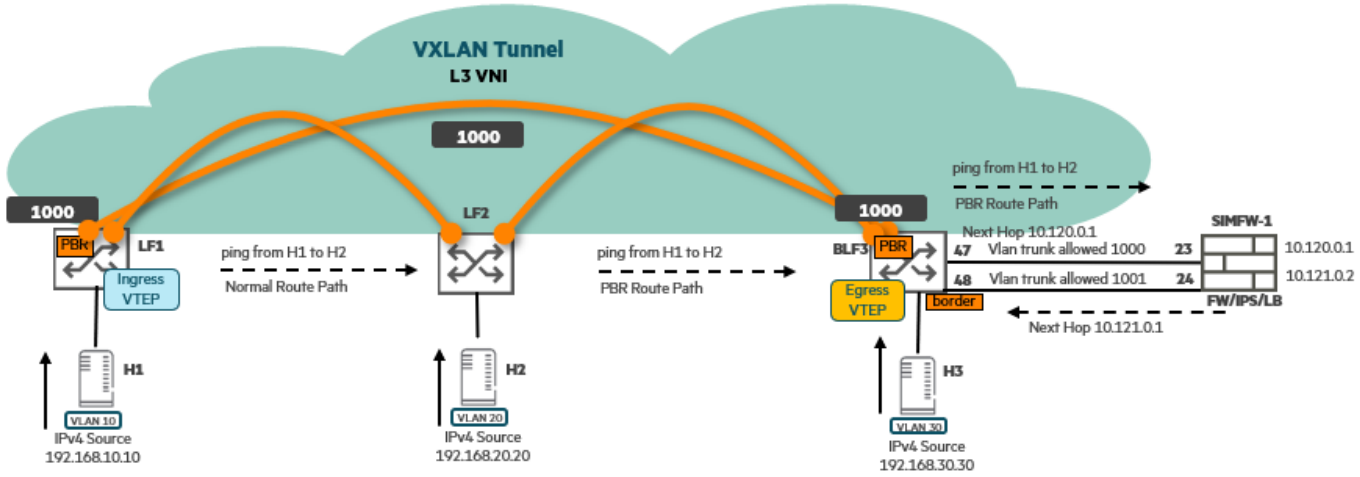

1. Check EVPN VXLAN PBR configurations

- Refer to configuration sections for IPv4/IPv6 sample configurations
- Check that your expected VRF, SVI, policy, class, action-list and Next Hop (active) appears as expected
- If not, check OSPF and BGP underlay and overlay connections are established between the leaf pairs and spines

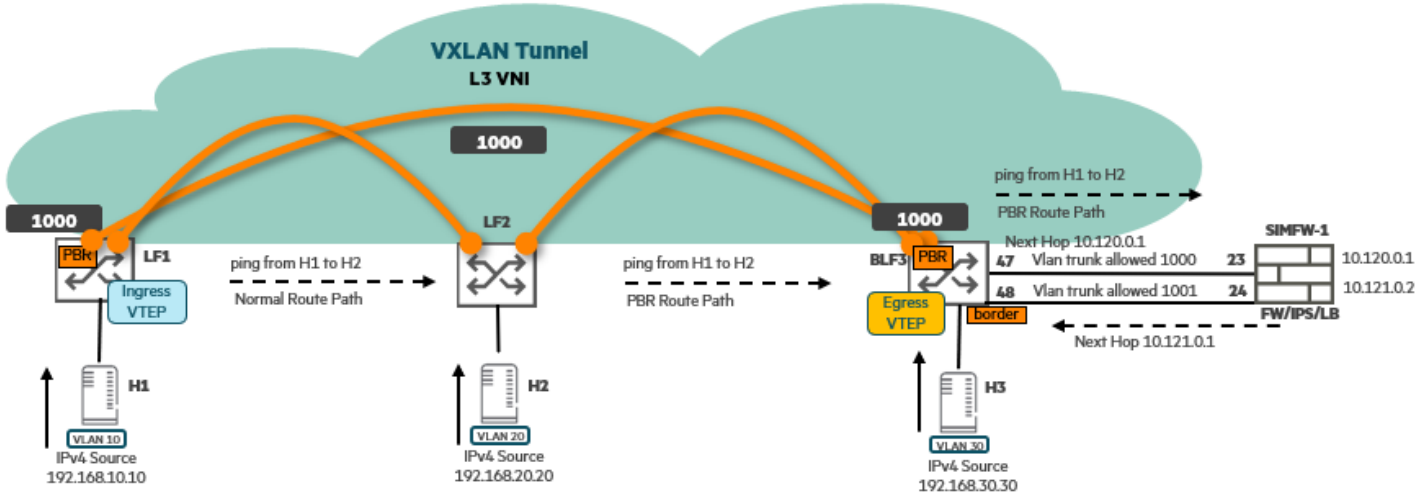


```
BLF3# sh pbr sum
VRF
      Port
                Policy
                          Class
                                    PBR
                                          Sequence  Type            Nexthop
-------------------------------------------------------------------------------
VRF1
      vni1000
                FW_Redirect
                          pbr_class
                                    pbr-nexthops
                                                10  nexthop         10.120.0.1 (active)
-------------------------------------------------------------------------
```

- Note: only active Next Hops are shown

2. Check policy hitcounts to verify traffic classes are matching

- Verity packet hitcounts on class matches are incrementing
- If hits are incrementing, indicated traffic flows will use the PBR policy

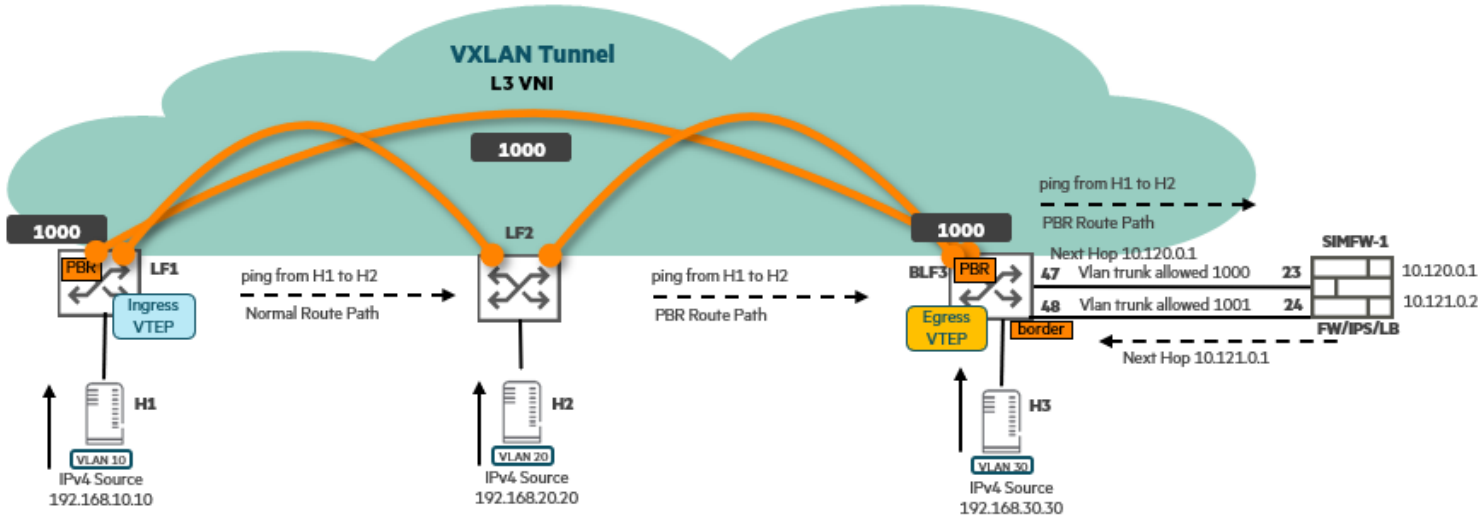


```
BLF3# show policy hitcounts FW_Redirect
```

```
BLF3# show policy hitcounts FW_Redirect
```

```
BLF3# show policy hitcounts FW_Redirect
Statistics for Policy FW_Redirect:

VRF VRF1
vni 1000 (routed-in):
     Matched Packets  Configuration
10 class ip pbr_class action pbr pbr-nexthops
                    40  10 match icmp 192.168.10.10 192.168.20.20 count
```

3. Check Next Hop Reachability via L2/L3 VNI

- For L2 Next Hop details, refer to the 10.09 release content
- L3 VNIs
  - Next Hop should be reached over L3VNI 1000 and **MUST** be of type **bgp_evpn**

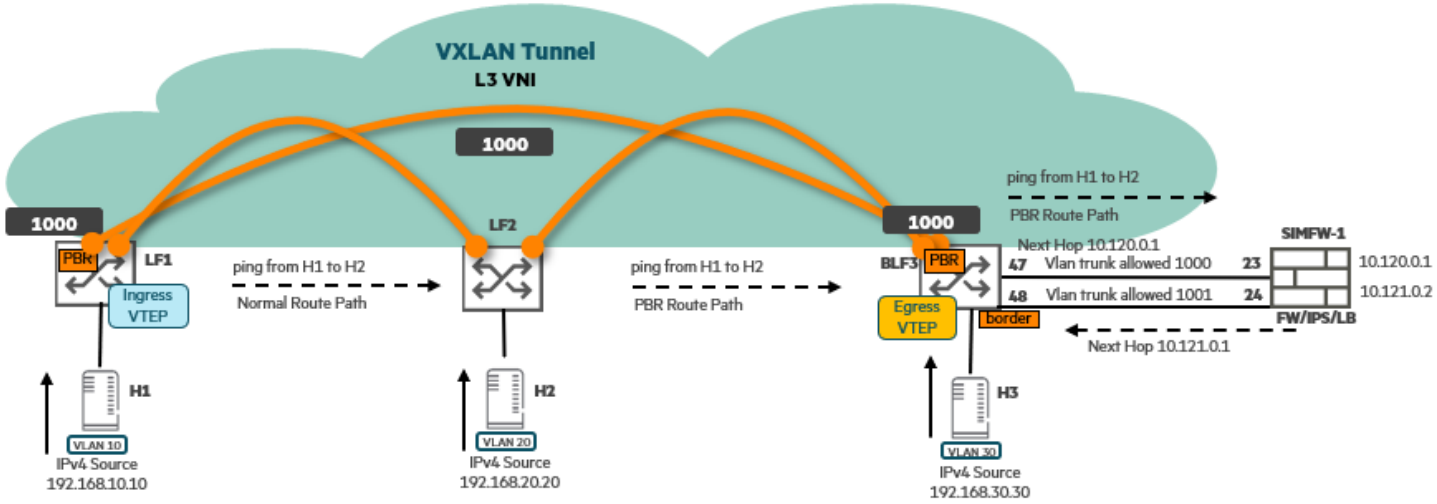


```
LF1# show ip route 192.168.20.0/24 vrf VRF1

VRF: VRF1

   Prefix           : 192.168.20.0/24                          VRF(egress)        : -
   Nexthop          : 2.2.2.2                                   Interface          : -
   Origin           : bgp                                       Type               : bgp_evpn
   Distance         : 200                                       Metric             : 0
   Age              : 05d:21h:11m                               Tag                : 0
   Encap Type       : vxlan                                     Encap Details      : l3vni 1000
```

```
LF1# ping 2.2.2.2 source 1.1.1.1
PING 2.2.2.2 (2.2.2.2) from 1.1.1.1 : 100(128) bytes of data.
108 bytes from 2.2.2.2: icmp_seq=1 ttl=63 time=0.244 ms
```

4. Verify traffic reaches the Next Hop IPv4/IPv6 address (capture)

- **When traffic is not received at the destination**, mirror ingress and egress traffic arriving and departing from the VTEP connected to the Next Hop; **BLF3** in this case
- **When traffic is not received at the Next Hop**, mirror ingress and egress traffic arriving and departing from the Ingress VTEP, i.e., **LF1**

```
BLF3# show mirror 1
Mirror Session: 1
Admin Status: enable
Operation Status: enabled
Source: vlan rx none
Source: vlan tx none
Source: interface 1/1/48 both
Destination: interface 1/1/3
Output Packets: 300
Output Bytes: 3000
```

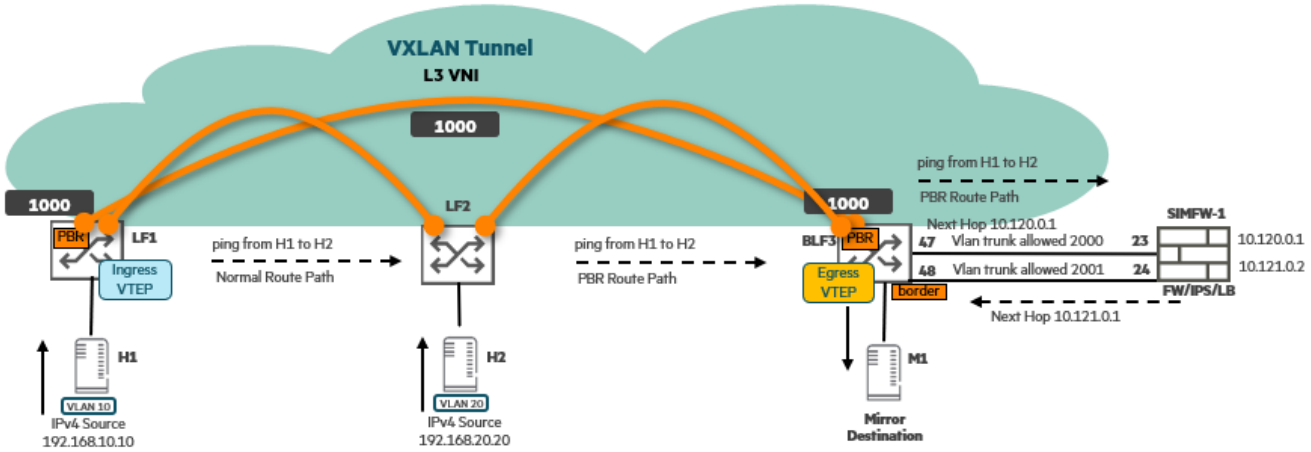

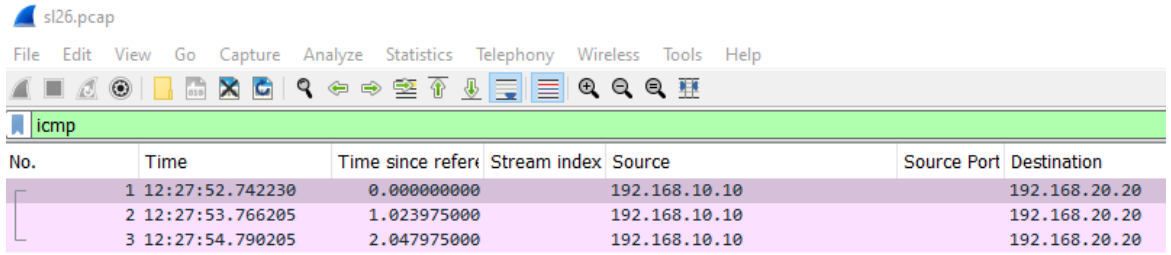

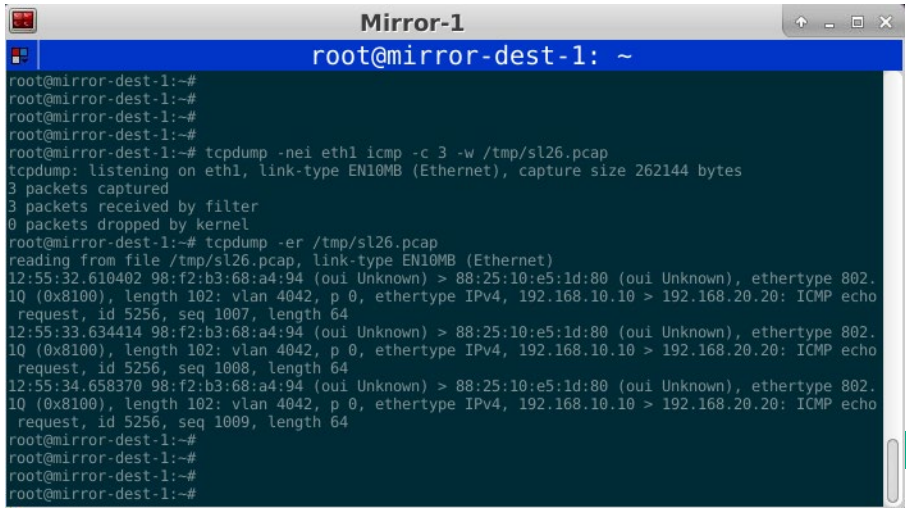

5. Check for L3 Forwarding Loops – TTL exceeded

- Traffic reaches the Next Hop IP address but never reaches the destination
- Ping **from 192.168.10.10** shows TTL exceeded
- In this example, if the PBR policy is also applied to **LF2**, then a route forwarding Loop will occur where return traffic from **BLF3** from the firewall is sent to **LF2**, which sends it back to **BLF3**
- **The PBR policy MUST be removed from LF2**

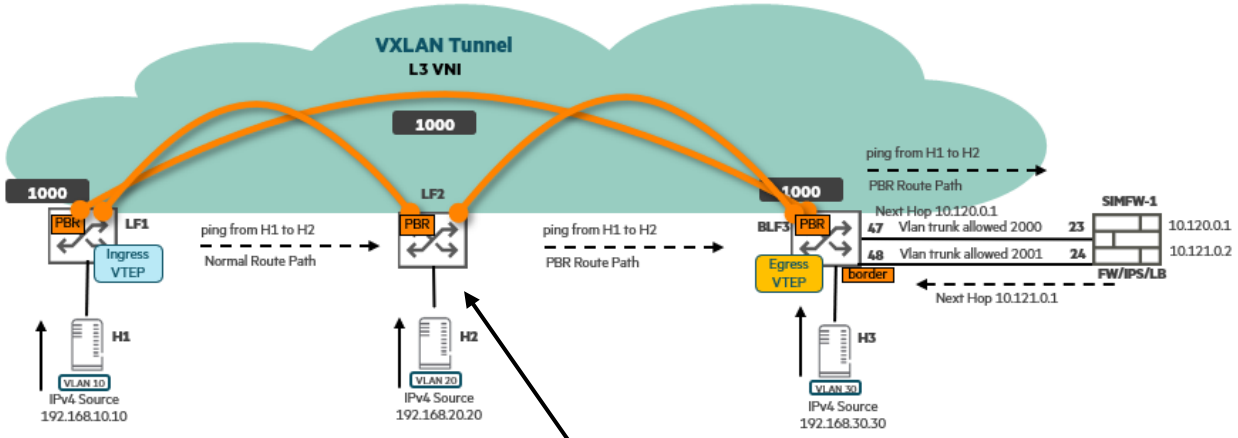


```
root@host1: ~
root@host1:~# ping 192.168.20.20
PING 192.168.20.20 (192.168.20.20) 56(84) bytes of data.
From 192.168.30.1 icmp_seq=1 Time to live exceeded
From 192.168.30.1 icmp_seq=2 Time to live exceeded
From 192.168.30.1 icmp_seq=3 Time to live exceeded
From 192.168.30.1 icmp_seq=4 Time to live exceeded
From 192.168.30.1 icmp_seq=5 Time to live exceeded
From 192.168.30.1 icmp_seq=6 Time to live exceeded
From 192.168.30.1 icmp_seq=7 Time to live exceeded
From 192.168.30.1 icmp_seq=8 Time to live exceeded
From 192.168.30.1 icmp_seq=9 Time to live exceeded
From 192.168.30.1 icmp_seq=10 Time to live exceeded
From 192.168.30.1 icmp_seq=11 Time to live exceeded
From 192.168.30.1 icmp_seq=12 Time to live exceeded
From 192.168.30.1 icmp_seq=13 Time to live exceeded
From 192.168.30.1 icmp_seq=14 Time to live exceeded
From 192.168.30.1 icmp_seq=15 Time to live exceeded
From 192.168.30.1 icmp_seq=16 Time to live exceeded
^C
--- 192.168.20.20 ping statistics ---
16 packets transmitted, 0 received, +16 errors, 100% packet loss, time 15016ms
```

PBR

```
class ip pbr-class
    10 match icmp 192.168.10.10 192.168.20.20
pbr-action-list pbr-nexthops
    10 nexthop 10.120.0.1
policy FW_Redirect
    10 class ip pbr-class action pbr pbr-nexthops

interface vxlan 1
   vni 1000
   apply policy FW_Redirect routed-in
```

6. Check resources for TCAM usage

- Show resources

```
leaf1# show resources

Resource Usage:

Mod  Description
         Resource                          Width    Used Reserved    Free
-------------------------------------------------------------------------
1/1  Routed Ingress IPv4 VLAN Policy
         Ingress TCAM Entries                  3    3432     6912
     Routed Ingress IPv6 VLAN Policy
         Ingress TCAM Entries                  3    3432        *
     Ingress Control Plane Policing
         Ingress TCAM Entries                  3     399     2304
     Egress Control Plane Policing
         Egress TCAM Entries                   2      92     1024
     Pre Ingress Control Plane Policing
     Total
         Ingress TCAM Entries                        7263     9216       0
         Egress TCAM Entries                           92     1024    1024
         Policers                                       0             6144
         Ingress L4 Port Ranges                         0               32
* This feature shares reserved resources with the preceding feature.
```

- Event Logs / Debug Messages
  - show events –s err –c tcam
  - show events –s err –c policy
  - show debug buffer module pbr
  - show debug buffer module classifier
  - show debug buffer module policy

```
leaf1# show events -s err -c tcam
-------------------------------------------------
Event logs from current boot
-------------------------------------------------
2023-10-20T01:00:29.920873+00:00 leaf1 ops-switchd[3226]: Event|10202|LOG_ERR|AMM|1/1|TCAM entry installation has failed in table Routed_Ingress_IPv6_VLAN_Policy
leaf1# show events -s err -c policy
-------------------------------------------------
Event logs from current boot
-------------------------------------------------
2023-10-20T01:00:29.924907+00:00 leaf1 ops-switchd[3226]: Event|10101|LOG_ERR|AMM|1/1|Policy p127 failed to apply on vlan227
```

Demo

EVPN VXLAN PBR Demo – Next Hop reached over L3VNI

- Traffic to 192.168.20.20 normally routes through the EVPN VXLAN Fabric to **LF2** as it is directly connected
- PBR policy applied to **VLAN interface 10** of ingress VTEP **LF1**
  - Required to alter the normal route path to the **PBR Next Hop 10.120.0.1**
- PBR policy applied to **L3VNI 1000** of Egress VTEP **BLF3** connected directly through to the firewall off SVI 1000
  - Return traffic from the simulated firewall is routed through **SVI 1001** back to **BLF3**
- Redirect ping traffic from 192.168.10.10 to destination 192.168.20.20 to Next Hop 10.120.0.1
- **Deny** ssh traffic from 192.168.10.10 to destination 192.168.20.20
- All other traffic to 192.168.20.20 continues to use **LF2** as the Next Hop in the VXLAN fabric

VRF	VLAN	SVI	Subnet	L2VNI	L3VNI
red	10	10	192.168.10.0/24	10	1000
	20	20	192.168.20.0/24	20	
	30	30	192.168.30.0/24	30	
	1000	1000	10.120.0.1/30	n/a	n/a
	1001	1001	10.121.0.1/30	n/a	n/a

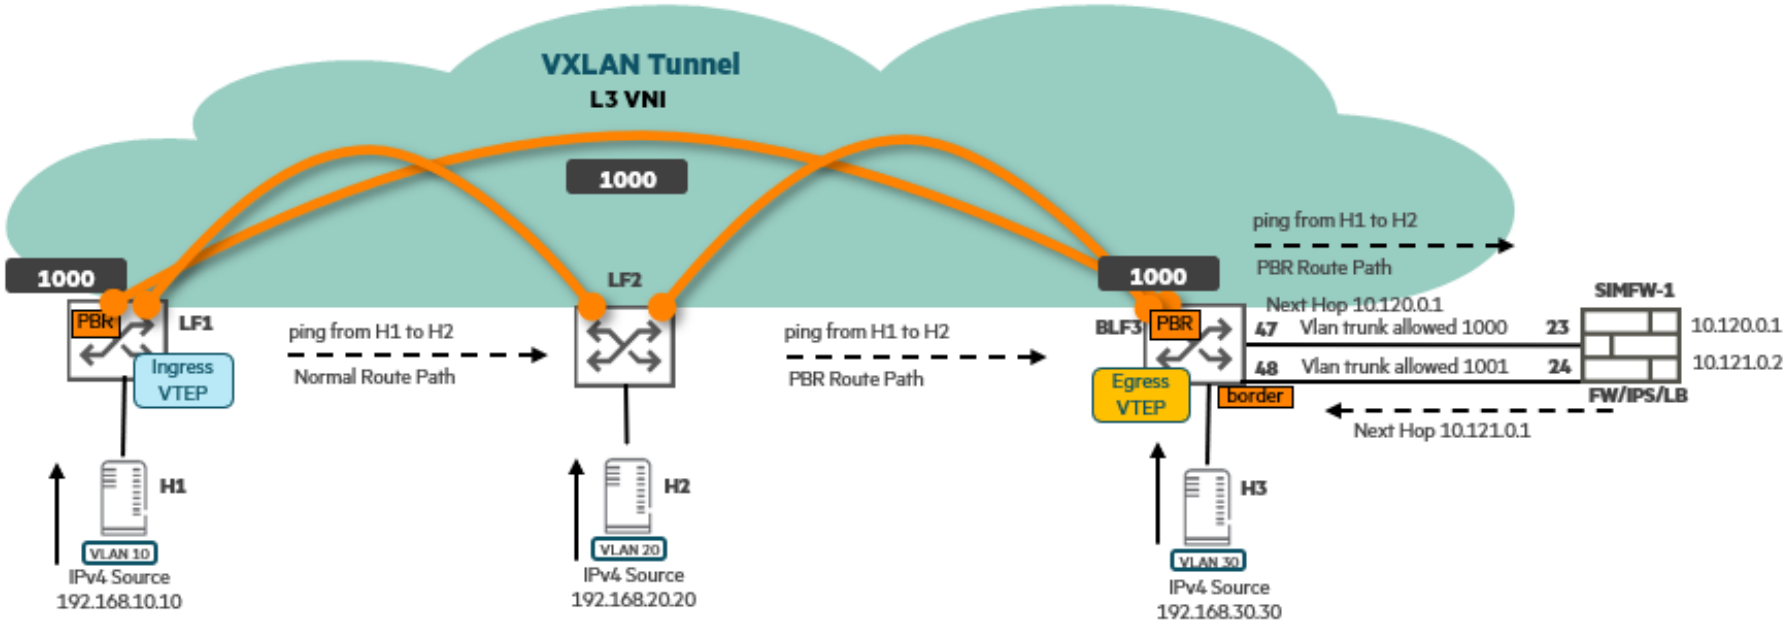

Thank you

aruba_switching_tme@hpe.com

© 2023 Hewlett Packard Enterprise Development LP