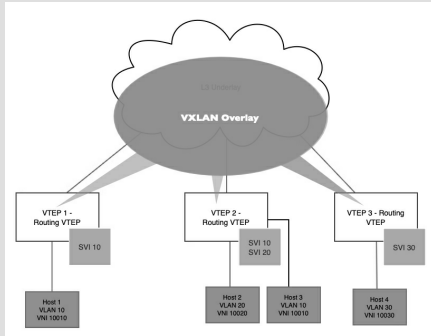


VXLAN Indirect Routing

To route between VLANs which are crossing the VXLAN overlay (also known as "stretched" or "extended" VLANs), Switched Virtual Interfaces (SVIs) need to be added to the VTEPs to perform the routing. With the indirect routing model, the dest host is not directly attached to the VTEP(s) where the SVIs are configured. As packets might need to go through multiple hops in the overlay to reach the final destination, this model is called "indirect". This scales better than direct routing as each VTEP is only responsible for routing a subset of the overlay subnets (fewer ARP entries, fewer routes.).

VXLAN and NAT cannot coexist on a switch.

Simple Topology



Troubleshooting

show interface vxlan 1

should be "up"; correctly reflect configured VLAN-to-VNI mappings; confirm control plane (multicast, HER, CVX, EVPN)

show mac address-table

The L2 forwarding table should show that mac addresses are either learned locally or from across the VXLAN overlay - "Vx1"; if we are not learning MACs from another VTEP confirm flood list and L3 reachability between VTEPs

show vxlan address-table

Shows the VXLAN MAC info, including the Host MAC, remote VTEP IP, and MAC moves.

show ip route

All VTEP IPs (primary and secondary) should have L3 reachability (ping to confirm)

show vxlan counters software

See "Arista VXLAN Bridging" for more info.

Troubleshooting (cont)

show platform fap arp

Confirm ARP in hardware.

show platform fap ip route

Confirms the prefix points to an L3FloodFap0 interface used for recirculation. These ports are 100G.

VXLAN encapsulated protocol packets (BGP/ISIS/OSPF) use the CppSystemL2Bcast queue.

Configuration

VTEP1

Configure a loopback as the L3 source interface for the VXLAN tunnel /VTEP.

!

interface loopback 1

```
ip address 1.1.1.1/32
```

Configure a second IP, the virtual VTEP IP, on routing VTEPs.

```
ip address 100.10.0.1 00.1 00/32 secondary
```

!

Add DGs for routing - aka "anycast gateway"; can use VARP as seen here or "ip address virtual" to save IP space.

!

interface vlan 10

```
ip address 10.0.0.1/24
```

```
ip virtual r outer address 10.0.0.254
```

!

Configure the VXLAN tunnel interface.

!

interface vxlan 1

Add the source interface configured earlier.

```
vxlan source -in interface loopback1
```

Configure the UDP destination port; be consistent across VTEPs.

```
vxlan udp-port 4789
```

Configure the VLAN to VNI bindings for any extended VLANs.

```
vxlan vlan 10 vni 10010
```

Configure the static flood-set if using HER for control plane; include all VTEP IPs that need to receive BUM traffic.

```
vxlan flood vtep 2.2.2.2 3.3.3.3
```

!

Configuration (cont)

```
> Configure the virtual MAC to be used with all virtual IPs.
!
ip virtual-router mac-address 00:00:11:22:33:44
!
Ensure routing is enabled.
ip routing
!
VTEP2
!
interface loopback 1
  ip address 2.2.2.2/32
Ensure the virtual VTEP IP is identical across all routing VTEPs.
  ip address 100.100.100.100/32 secondary
!
interface vlan 10
  Here using "ip address virtual" as the anycast gateway IP.
  ip address virtual 10.0.0.254/24
!
interface vlan 20
  ip address virtual 20.0.0.254/24
!
interface vxlan 1
  vxlan source-interface loopback1
  vxlan udp-port 4789
  vxlan vlan 10 vni 10010
  vxlan vlan 20 vni 10020
  vxlan flood vtep 1.1.1.1 3.3.3.3
!
ip virtual-router mac-address 00:00:11:22:33:44
!
ip routing
!
VTEP 3
!
interface loopback 1
```

Configuration (cont)

```
> ip address 3.3.3.3/32
  ip address 100.100.100.100/32 secondary
!
interface vxlan 1
  vxlan source-interface loopback1
  vxlan udp-port 4789
  vxlan vlan 30 vni 10030
  vxlan flood vtep 1.1.1.1 2.2.2.2
!
ip routing
```

All routing VTEPs have secondary IPs, called Virtual VTEP IPs, configured that are identical.

In order to learn routes to remote destinations, routing protocols such as OSPF/BGP/ISIS can be configured in the overlay.

There is no "indirect" specific configuration necessary.

