

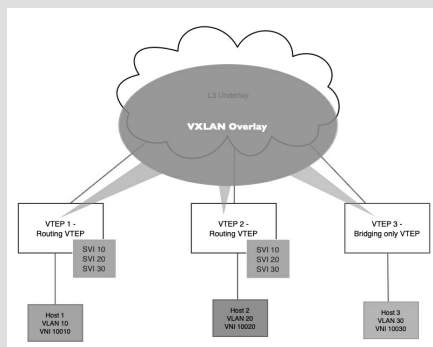
### VXLAN Direct Routing

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 direct routing model, every VTEP serves as the first-hop default gateway (DG) for locally attached hosts. Every SVI will be configured on every VTEP, even if there isn't a local host for a particular extended VLAN.

<https://www.arista.com/en/um-eos/eos-section-22-3-vxlan-configuration#ww1154798>

### VXLAN and NAT cannot coexist on a switch.

### Simple Topology



This network includes 3 VTEPs - 2 Routing VTEPs and 1 Bridging-only VTEP. This example is used to illustrate how and why the virtual VTEP IP (secondary IP on the VTI loopback) is used.

### Pros of Direct Routing

Leads to more efficient traffic flows than indirect routing.

Load sharing and redundancy as there will be more than one VTEP who can route .

Useful in vMotion as all VTEPs have the same VIP and VMAC so default gateway configurations remain the same for servers across Data Centers.

### Troubleshooting

#### show int 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*

#### show vxlan counters varp

*Gives insight into ARP replies/requests sent and received, etc.*

### With MLAG, make sure loopback primary IPs are identical.

Make sure there is a secondary IP configured on the loopback for routing VTEPs.

### Run a Recirc/TCAM profile check based on platform

*See Arista VXLAN Routing*

Confirm MTU on uplinks

*VXLAN adds a 50-byte header to the packet.*

### Confirm the VMAC on any bridging VTEPs is either not configured or is a different VMAC than the routing VTEPs.

There should be no secondary IP on the bridging VTEP's loopback.

### Secondary IP should be in the flood list of any bridging VTEPs.

Make sure SVIs are configured on EVERY routing VTEP when using the Direct Routing model.



### 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 100.100.100.100/32 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-router address 10.0.0.254
!
interface vlan 20
  ip address 20.0.0.1/24
  ip virtual-router 20.0.0.254
!
interface vlan 30
  ip address 30.0.0.1/24
  ip virtual-router address 30.0.0.254
!
Configure the VXLAN tunnel interface
!
interface vxlan 1
Add the source interface configured earlier
  vxlan source-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
```

### Configuration (cont)

```
vxlan vlan 20 vni 10020
vxlan vlan 30 vni 10030
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
!
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 vlan 30
  ip address virtual 30.0.0.254/24
!
interface vxlan 1
  vxlan source-interface loopback1
```



### Configuration (cont)

```

vxlan udp-port 4789
vxlan vlan 10 vni 10010
vxlan vlan 20 vni 10020
vxlan vlan 30 vni 10030
vxlan flood vtep 1.1.1.1
!
ip virtual-router mac-address 00:00:11:22:33:44
!
ip routing
!
VTEP 3
!
interface loopback 1
  ip address 3.3.3.3/32
No virtual VTEP IP on a bridging VTEP!
!
interface vxlan 1
  vxlan source-interface loopback1
  vxlan udp-port 4789
  vxlan vlan 10 vni 10010
  vxlan vlan 20 vni 10020
  vxlan vlan 30 vni 10030
Ensure the virtual VTEP IP is included in the flood-set for BUM traffic.
  vxlan flood vtep 1.1.1.1 2.2.2.2 100.100.100.1-00
!
ip routing

```

VTEP 1 and 2 have secondary IPs configured that are identical. This is to avoid MAC flapping on the bridging VTEP 3. If the secondary IP wasn't configured, when VTEP 1 and 2 send GARPs out for the SVIs using VARP, VTEP 3 would see the virtual MAC flap between VTEP 1 and VTEP 2 in the vxlan address table.

"ip address virtual" could also have been used, saving IP address space as no physical IP address is needed.

### Virtual MAC on the Bridging VTEP...

Notice there is no VMAC configured on the bridging VTEP. It is not unlikely that a VTEP could be a "logical" VTEP - i.e. two VTEPs in MLAG and have a VMAC configured for VARP. It is not unlikely that there are both regular VLANs, not using VXLAN at all, and VXLAN VLANs. If the VXLAN VLANs had no SVIs, making this a "bridging" VTEP, make sure the VMAC on these switches is different from the routing VTEPs for the VXLAN VLANs or the switches will consume the packets.

