


VXLAN Routing

VXLAN routing is the routing of traffic between extended L2 domains using VTEPs as the default gateways. Routing can be achieved via a direct or indirect routing model. In the "direct" version, packets bridge to their gateway, route on that same device, then bridge across the VXLAN overlay. With the "indirect" model, only a subset of the overlay tenant networks will get locally routed instead of all of the overlay networks. The other networks will use a transit VNI for routing.

If needed, see "Arista VXLAN Bridging Cheat Sheet" for a review of the basics of VXLAN.

Example Topology



VXLAN and NAT cannot coexist on a switch.

VXLAN and MPLS cannot coexist on a switch.

VXLAN Routing by Platform

Platform	Additional config needed due to asic limitations (as of 4.18.1F)
7150	Good to go
7050X	Configure recirculation
7050X2	Good to go
7050X3,7300 X3	Good to go
7250X	Good to go
7060CX, 7260QX	Configure recirculation
7280E	Configure TCAM profile using "hardware tcam profile vxlan-routing"
7500E	Not supported
7280R,- 7500R	Configure TCAM profile using "hardware tcam profile vxlan-routing"

Please see block on how to configure recirculation channels.

Recirculation using front-panel ports

```
interface Recirc-Channel627
    switchport recirculation features vxlan
!
interface Ethernet1/2
    traffic-loopback source system device mac
    channel-group recirculation 627
!
interface Ethernet1/3
    traffic-loopback source system device mac
    channel-group recirculation 627
```

For more information:

<https://www.arista.com/assets/data/pdf/user-manual/um-eos/Chapters/VXLAN.pdf>

Recirculation using UnconnectedEthernet Interfaces

```
service interface unconnected expose
!
interface UnconnectedEthernet2
    traffic-loopback source system device mac
    channel-group recirculation 627
!
```

On the 7050X (but not on the 7060CX or 7260QX series), the entire chip BW is not used by the front panel ports and can be used for recirculation via this method. It is recommended to set the recirculation BW to match the uplink BW facing the spines. To learn more about this, please contact Arista or see the below:

<https://www.arista.com/assets/data/pdf/user-manual/um-eos/Chapters/VXLAN.pdf>

Confirmation

show recirc-channel X To confirm configuration

Base Configuration of a Routing VTEP

```
interface Vxlan1
    vxlan source-interface Loopback1
    vxlan udp-port 4789
    vxlan vlan 15 vni 10015
    vxlan vlan 240 vni 8358534
    vxlan vlan X vni Y
    vxlan flood vtep <remote VTEP IP> <remote VTEP IP>
!
```

Base Configuration of a Routing VTEP (cont)

```
interface Loopback1
  ip address 1.0.1.1/32 << VTEP IP
  ip address 1.0.1.2/32 secondary << Virtual VTEP
  IP (VVTEP)
!
ip virtual-router mac-address 00:00:11:22:33:44
!
ip routing
!
```

Notice the only difference with this portion of the config is adding a secondary IP known as the Virtual VTEP IP. This is only needed on routing VTEPs and not on bridging-only VTEPs. On bridging VTEPs, this VVTEP IP will be included in the flood list to ensure BUM traffic reaches all VTEPs. Don't forget to add the virtual MAC which binds with all Virtual IPs configured on the switch. See "Arista VXLAN Direct Routing" for additional information.

Direct or Indirect Routing?

See "Arista VXLAN Direct Routing" or "Arista VXLAN Indirect Routing" Cheat Sheets for details.

Configuration of the SVIs (Anycast Gateways)

Using VARP:	Using "ip address virtual":
interface Vlan15	interface Vlan15
ip address 1.0.15.1/24	ip address virtual 1.0.15.254/24
ip virtual-router address 1.0.15.254	
ip virtual-router mac-address 00:00:11:22:33:44	ip virtual-router mac-address 00:00:11:22:33:44

There are two options when configuring anycast gateways. VARP, commonly seen with MLAG, uses more IP addresses - a physical and a virtual IP address. "ip address virtual" only uses a virtual IP address. There are some considerations when using "ip address virtual" - see the relevant block for more info. For more info on VARP, please see the "Arista VARP" Cheat Sheet.

Considerations with "ip address virtual"

Must configure a secondary IP (matching MLAG peer and all other routing VTEPs) on VXLAN loopback for virtual VTEP IP.

Routing adjacencies can't form over SVIs configured with "ip address virtual".

In MLAG config, ARP responses to the SVI's VIP are synced with the MLAG peer to ensure consistency between the peers. This is achieved via the VXLAN agent so the "ip address virtual" option is only available in a VXLAN config.

No periodic GARP is sent (like is sent in VARP).

For non-VXLAN VLANs, use VARP (ip virtual-router address).

Configuring a TCAM hardware profile post 4.20.1F

```
sw(config)#hardware tcam
sw(config-hw-tcam)#system profile vxlan-routing
sw(config-hw-tcam)#show hardware tcam profile
```

ARP and MAC Aging Timers

To reduce the potential for traffic flooding when the MAC ages out, configure (interface level) the ARP aging timer to equal or less than the MAC Aging timer.

```
(config-if)#arp timeout 14400
(config)#mac address-table aging-time 14400
```

