Traffic Load Balancing in EVPN/VXLAN Networks

Tech Note

December 2017

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Contents

Introduction

This document focuses on how traffic is load balanced in an EVPN/VXLAN network using QFX Series devices.

An EVPN/VXLAN network consists of a physical underlay layer and the logical overlay layer. This document illustrates how traffic is load balanced at both layers using QFX10000 Series switches, which use Juniper's custom PE chip, and QFX5100/QFX5110 switches, which use Broadcom Trident2/Trident2+ chips. It also highlights the differences between these platforms when it comes to selecting next hops, and how these differences affect the platforms' load-balancing capabilities. Traffic scenarios that will be covered include:

- Intra-VNI traffic using EVPN type-2 routes
- Inter-VNI traffic using EVPN type-2 routes
- Inter-VNI traffic using EVPN type-5 routes

For more information on VXLAN network identifiers (VNIs), see Understanding EVPN with VXLAN Data Plane Encapsulation. For more information on EVPN route types, see Understanding EVPN Pure Route Type-5 on QFX Series Switches.

Topologies

The topologies used to demonstrate traffic flows using EVPN type-2 and type-5 routes are shown in Figures 1 and 2.

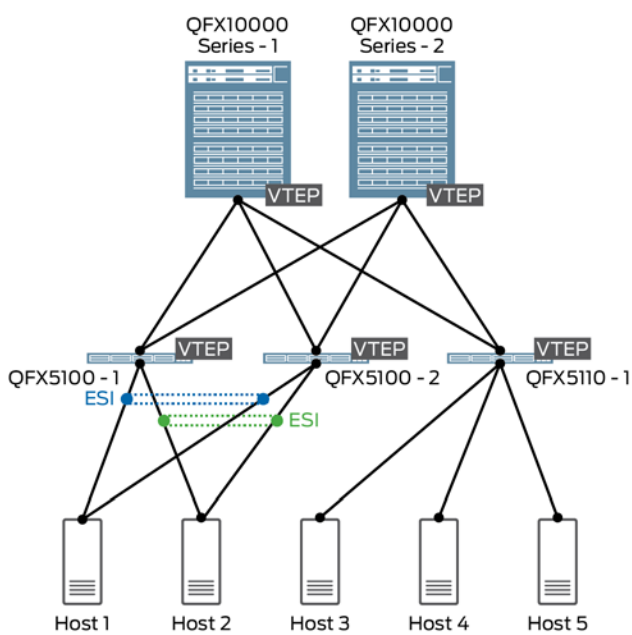

Figure 1: Topology for traffic flows using EVPN type-2 routes

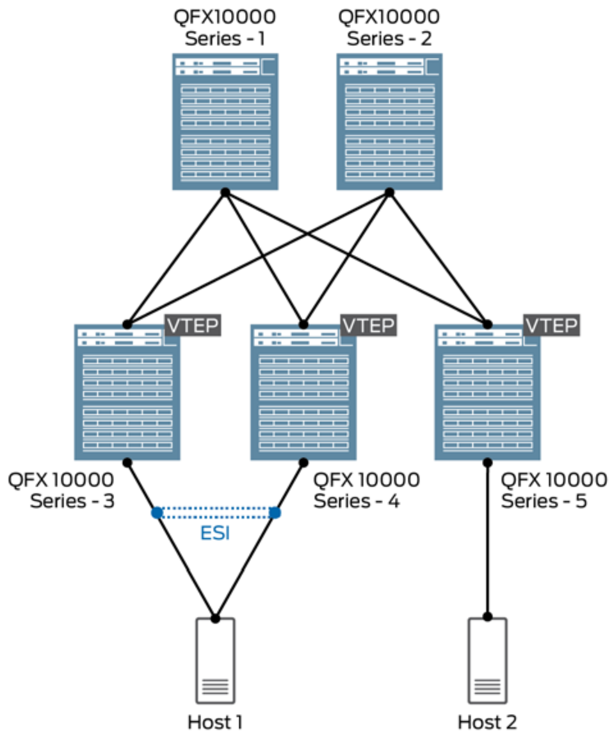

Figure 2: Topology for traffic flows using EVPN type-5 routes

Technology Overview

Before getting into some load balancing scenarios, it is important to understand some key concepts.

Underlay Network

In an EVPN/VXLAN environment, the underlay network is the fabric that provides physical connectivity between all the devices in the network. Also known as a Clos network, the underlay typically uses a leaf-spine design with Layer 3 routing providing reachability between the devices. The goal of this network is to provide any-to-any connectivity between all devices, as well as to provide inter-device reachability for the signaling protocols used in the overlay network.

In this document, EBGP is used in the underlay network, with each device having its own ASN. No IGP is required. An example of this setup is shown in Figure 3.

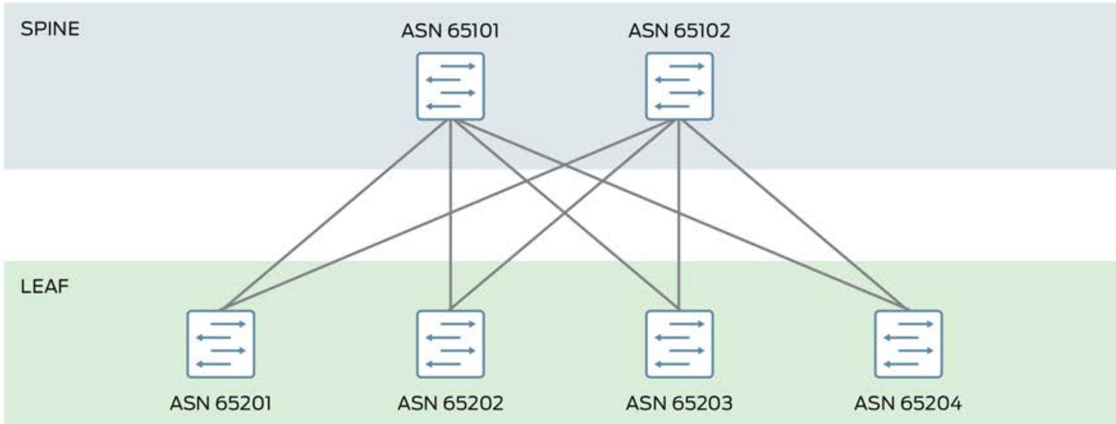

Figure 3: Leaf-spine design; ASNs assigned per device

Load Balancing in the Underlay

To ensure packets are load balanced in the underlay, the `multipath multiple-as` parameter must be added on all devices, since by default only prefixes advertised by neighbors in the same AS are considered. The full configuration statement is as follows:

```
set protocols bgp group underlay multipath multiple-as
```

A load balancing policy must also be defined and applied to the forwarding table, as follows:

```
set policy-options policy-statement LB then load-balance per-packet

set routing-options forwarding-table export LB
```

Overlay Network

Network overlays are created by encapsulating traffic and tunneling it over a physical network. A number of tunneling protocols can be used in the data center to create network overlays—the most common protocol is Virtual Extensible LAN (VXLAN). The VXLAN tunneling protocol encapsulates Layer 2 Ethernet frames in Layer 3 UDP packets to enable virtual Layer 2 subnets or segments that can span the underlying (physical) Layer 3 network.

In a VXLAN overlay network, each Layer 2 subnet or segment is uniquely identified by a virtual network identifier (VNI). A VNI enables segmenting of traffic the same way that a VLAN ID segments traffic. As is the case with VLANs, endpoints with the same VNI can communicate directly with each other, whereas endpoints on different VNIs require a router, or gateway.

The entity that performs VXLAN encapsulation and decapsulation is called a VXLAN tunnel endpoint (VTEP). VTEPs typically reside in hypervisor hosts, such as ESXi or KVM hosts, but can also reside in network devices to support bare-metal server (BMS) endpoints. Each VTEP is typically assigned a unique IP address.

Though VXLANs can be manually provisioned, typically a signaling protocol is used. Ethernet VPN (EVPN) is a standards-based protocol that provides virtual multipoint bridged connectivity between different domains over an IP or IP/MPLS backbone network. This control-plane technology uses Multiprotocol BGP (MP-BGP) for MAC and IP address (endpoint) distribution, with MAC addresses being treated as "routes." As used in data center environments, EVPN enables devices acting as VTEPs to exchange reachability information with each other about their endpoints.

In this document, MP-IBGP peering is used, with the EVPN protocol family (`family evpn`) enabled. When using a 'centralized routing' model, i.e. Layer 3 gateways on the spine devices and Layer 2 gateways on the leaf devices, MP-IBGP sessions are established between all nodes; when using the 'distributed routing', i.e. a collapsed Layer 2 / Layer 3 gateway on the leaf devices, MP-IBGP sessions are needed only between leaf devices. Route reflection can be used to reduce the number of peering sessions.

Why is Load Balancing Needed in the Overlay?

Server multihoming to redundant top-of-rack devices is a common requirement in data centers. Traditionally, this requirement required proprietary solutions such as multichassis link aggregation (MLAG), multichassis link aggregation groups (MC-LAGs), Virtual Chassis Port (VCP), switch stacking, and Virtual Chassis. While each solution has its merits, each requires use of a single vendor's devices. And from a technical standpoint, when using MLAG/MC-LAG multihoming is limited to two PE devices.

The standards-based EVPN protocol, on the other hand, includes built-in multihoming capabilities, scales horizontally across any number of PE devices, and seamlessly integrates into multivendor, Layer 3 Clos fabrics.

ESIs and Multihoming

In Figure 4, H2 is multihomed via a standard LAG to both LS2 and LS3 in the same Layer 2 domain. On these leaf switches, this common Layer 2 domain takes the form of an Ethernet segment (ES), with a common Ethernet Segment Identifier (ESI) assigned. Both LS2 and LS3 advertise direct reachability to this segment via a Type1 route to LS1.

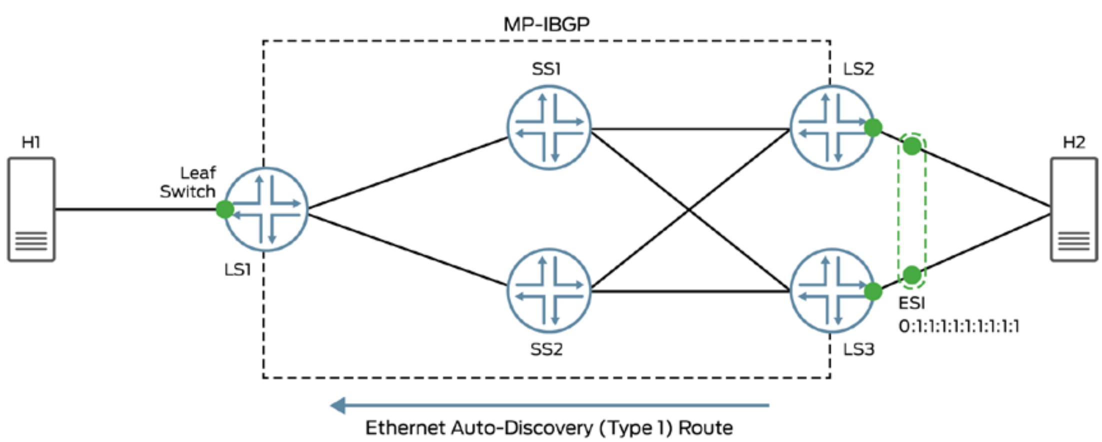

Figure 4: ESI advertisement via EVPN Type 1 route

Type 1 routes do not advertise MAC addresses for endpoints learned on this ESI. For MAC reachability, a Type 2 route is required. For the moment, let's assume that LS2 and LS3 have both learned H2's MAC address.

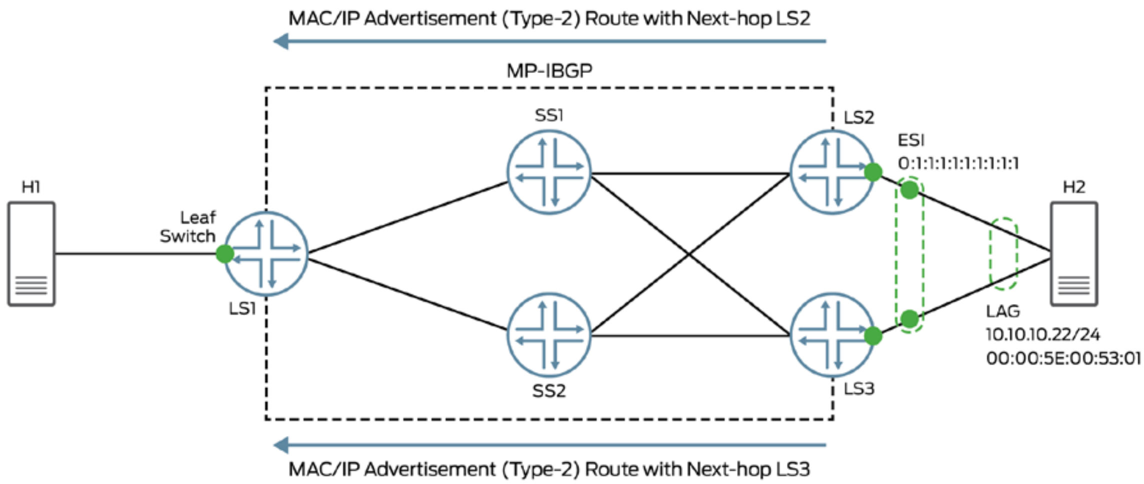

Figure 5: EVPN Type 2 advertisement with associated ESI

In Figure 5, LS1 receives a Type 2 advertisement with H2's MAC address from LS3, with associated ESI 0:1:1:1:1:1:1:1:1:1; it similarly receives a Type 2 advertisement for H2 from LS2 with the same ESI. Given this information, LS1 knows that H2 is reachable via both peers.

With multiple paths established, LS1 will load balance traffic destined to H2 through the VXLAN tunnels to both LS2 and LS3, as shown in Figure 6.

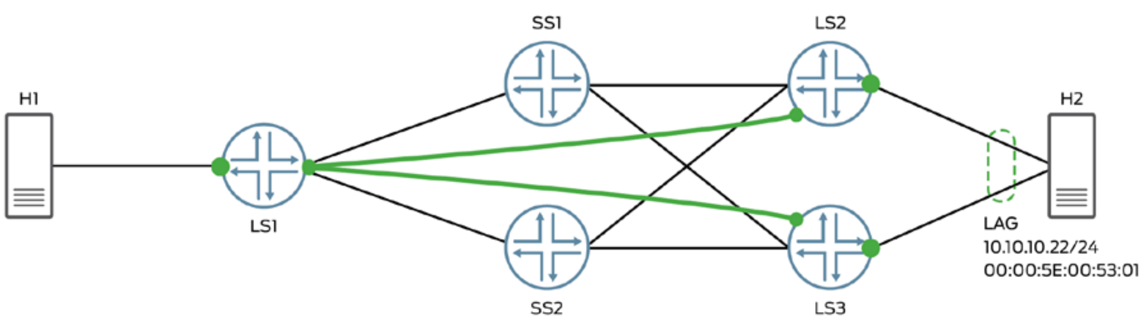

Figure 6: Multipathing from LS1 to H2 via LS2 and LS3

Aliasing

A problem arises, however, when only one of LS2 and LS3 has learned H2's MAC address, as shown in Figure 7.

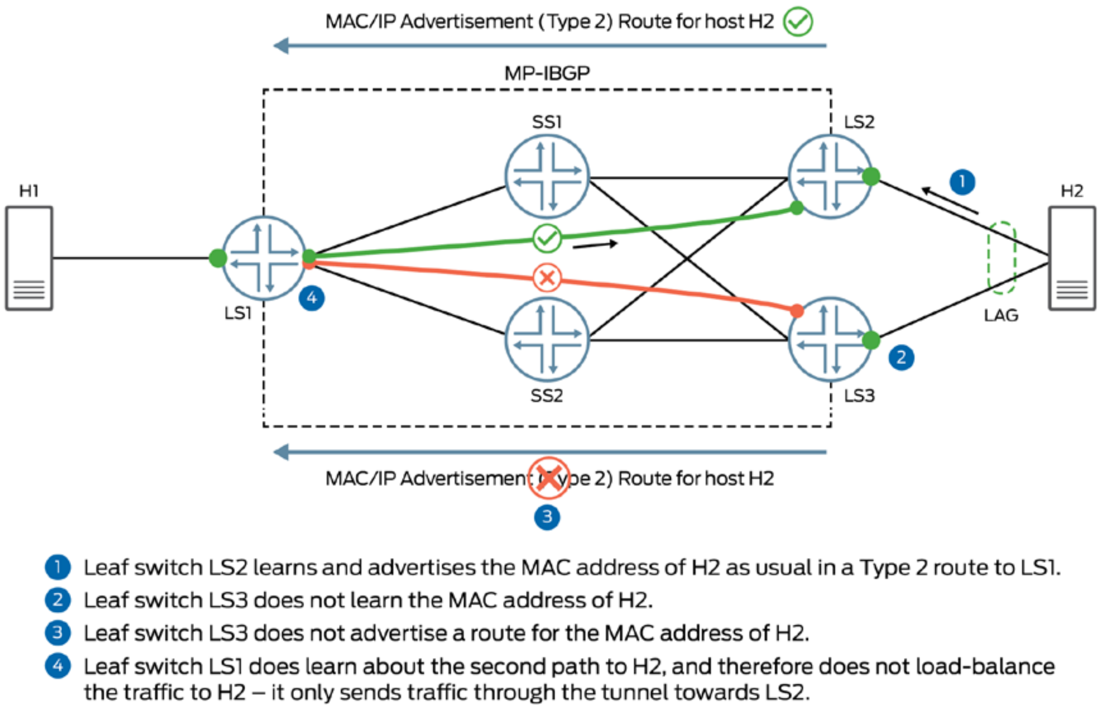

Figure 7: Multipathing failure from LS1 to H2 via LS2 and LS3

The solution to this problem is aliasing. Aliasing is the ability of a remote leaf switch (in this case LS1) to load balance Layer 2 unicast traffic towards a given end endpoint (H2) through all leaf switches connected to the endpoint's Ethernet segment (LS2 and LS3), even when the remote leaf switch has not received an endpoint MAC address advertisement from all those leaf switches.

With aliasing, the multihoming issue in Figure 7 is overcome, as shown in Figure 8.

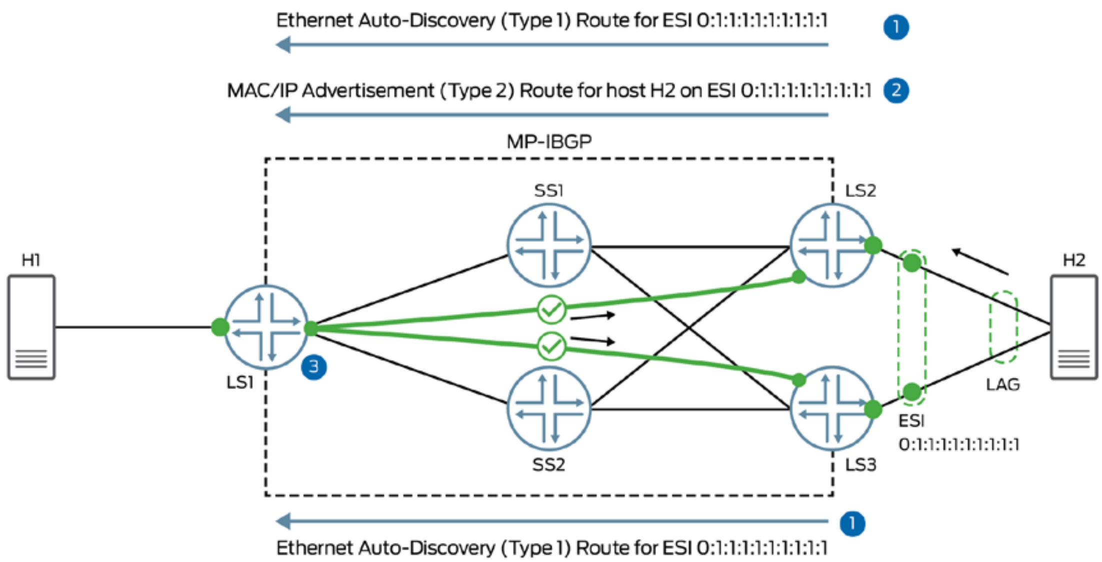

1. LS1 receives the Type 1 route from both LS2 and LS3, as usual.
2. LS1 receives a Type 2 route only from LS2.
3. LS1 knows on which ES host H2 is located.
   LS1 knows that since both LS2 and LS3 have advertised the same ESI, H2 must be reachable via both leaf switches.
   LS1 can load-balance traffic to host H2 over both tunnels.

Figure 8: Multipathing from LS1 to H2 via LS2 and LS3, with aliasing

Load Balancing: Underlay vs. Overlay

As discussed above, a set of leaf switches multihomed to a local end host will each advertise reachability for that host, resulting in the remote leaf switch(es) learning multiple paths to reach the end host. There are multiple load-balancing considerations here. In order to reach a particular (i.e. single) remote VTEP, traffic can flow through multiple paths via different spine nodes; this is taken care of using standard load balancing in the underlay. However, using all the available paths (i.e. all the remote VTEPs) to the end host in an active-active manner requires using load-balancing capabilities in the overlay context.

The rest of this document focuses on how load balancing can be achieved in the overlay, and identifies some of the differences in load balancing capabilities across QFX Series platforms.

Load Balancing Scenarios

Scenario 1a: Intra-VNI traffic using EVPN Type-2 routes – multiple flows to same DMAC

Topology

- Layer 3 gateway at spine layer (QFX10000)
- Layer 2 gateway at leaf layer (QFX5100 / QFX5110)

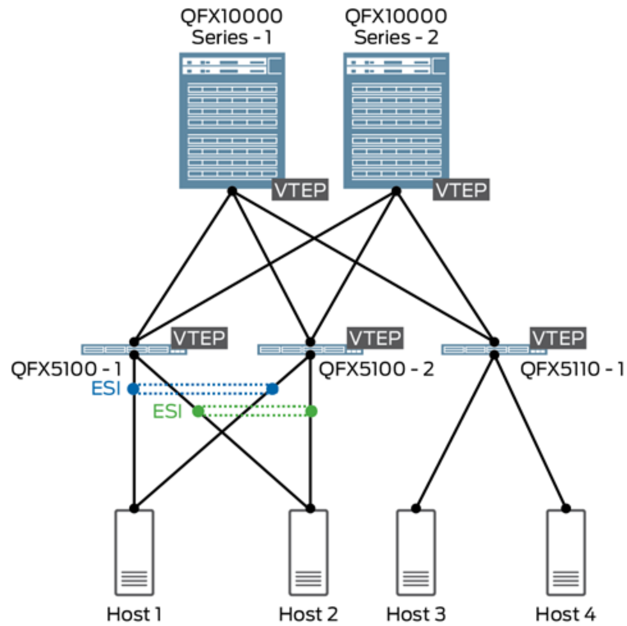

Traffic Flow

- Sending 50 flows from H3 (10.10.10.110) to H1 (10.10.10.10)

Traffic details:

- SMAC: 00:10:94:00:00:23
- DMAC: 00:10:94:00:00:24
- SIP: 10.10.10.110
- DIP: 10.10.10.10
- SPORT: 1024-1074
- DPORT: 50000 - 50050

Operation

QFX5110-1 has received H1's MAC address via either QFX5100-1 or QFX5100-2. As discussed earlier, due to aliasing, QFX5110-1 knows that it can reach H1's MAC address via both QFX5100-1 and QFX5100-2, even though only one of them sent the Type 2 advertisement.

The QFX5100/QFX5110 can only install VTEP next hops in the PFE; it cannot install ESI next hops. This means that, for any given overlay destination, only one remote VTEP can be selected. To send traffic to the selected VTEP, traffic can be load balanced at the underlay layer through the two spine nodes.

Hence, for a given destination MAC address, there is no load balancing of the flows at the overlay layer; traffic is load balanced only in the underlay.

Verification

SMAC (H3): 00:10:94:00:00:23
DMAC (H1): 00:10:94:00:00:24

Check the MAC table on QFX5110-1 for H3 and H1:

```
{master:0}
user@5110-1> show ethernet-switching table | match 00:23
   bd10                00:10:94:00:00:23   D        xe-0/0/46.0

{master:0}
user@5110-1> show ethernet-switching table | match 00:24
   bd10                00:10:94:00:00:24   DR
00:01:01:01:01:01:01:01:01:01
```

The output shows that QFX5110-1 has learned H3's MAC address through a local interface, and H1's MAC address over ESI 00:01:01:01:01:01:01:01:01:01.

Verify which remote VTEPs that are advertising ESI 00:01:01:01:01:01:01:01:01:01 to QFX5110-1:

```
user@5110-1> show ethernet-switching vxlan-tunnel-end-point esi | find 1752
00:01:01:01:01:01:01:01:01:01 default-switch           1752  131076  esi.1752
2
    RVTEP-IP             RVTEP-IFL      VENH     MASK-ID   FLAGS
    172.23.0.1            vtep.32772     1753     1          2
    172.24.0.1            vtep.32771     1751     0          2
```

The output shows the remote VTEPs and VTEP IFLs related to the ESI.

Check reachability information for the remote VTEPs:

```
user@5110-1> show route forwarding-table destination 172.23.0.1
Routing table: default.inet
Internet:
Enabled protocols: Bridging,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
172.23.0.1/32         user     1                 ulst   131070    20
                              192.168.90.1       ucst     1740    13 et-0/0/49.0
                              192.168.100.1      ucst     1741    13 et-0/0/50.0

user@5110-1> show route forwarding-table destination 172.24.0.1
Routing table: default.inet
Internet:
Enabled protocols: Bridging,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
172.24.0.1/32         user     1                 ulst   131070    20
                              192.168.90.1       ucst     1740    13 et-0/0/49.0
                              192.168.100.1      ucst     1741    13 et-0/0/50.0
```

The output shows ECMP next hops for both remote VTEPs.

Now login to the PFE to see the next hop installed for H1's MAC address:

```
FPC0(5110-1 vty)# show l2 manager mac-table

...
mac address        BD     learn  Entry  entry       hal         hardware info
                   Index  vlan   Flags  ifl         ifl         pfe  mask  ifl
  ---------------------------------------------------------------------------------------
  00:00:5e:00:01:01  3      0    0x0014 vtep.32770 vtep.32770   0    0x1  vtep.32770
  00:10:94:00:00:23  3      0    0x0814 xe-0/0/46.0 xe-0/0/46.0 0    0x1  xe-0/0/46.0
  00:10:94:00:00:24  3      0    0x0014 vtep.32771 vtep.32771   0    0x1  vtep.32771
...
```

Then verify the remote VTEP's IP address:

```
user@5110-1> show interfaces vtep.32771 | match endpoint
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 172.24.0.1, L2 Routing Instance:
default-switch, L3 Routing Instance: default
```

The output shows that the device has installed the VTEP next hop, and not the ESI next hop. For H1's MAC address, the VTEP installed is 172.24.0.1.

Scenario 1b: Intra-VNI traffic using EVPN Type-2 routes – multiple flows to multiple DMACs

Topology (same as Scenario 1a)

- Layer 3 gateway at spine layer (QFX10000)
- Layer 2 gateway at leaf layer (QFX5100 / QFX5110)

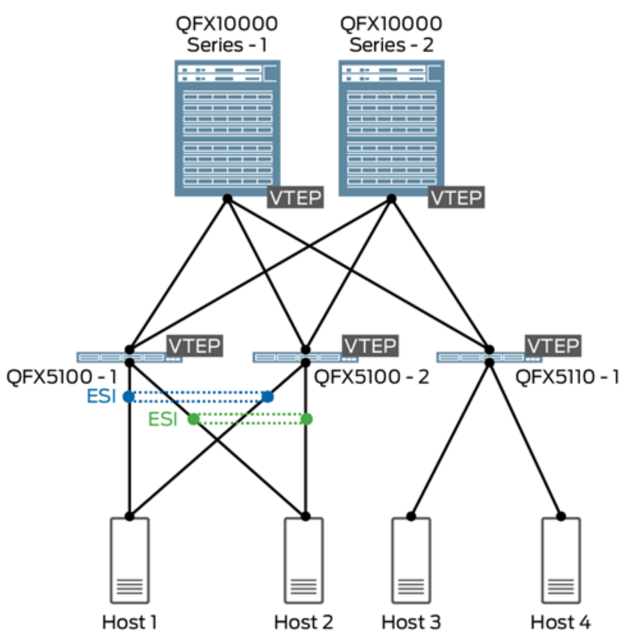

Traffic Flow

- Sending 50 flows from H3 (10.10.10.110) to H1 (10.10.10.10)
- Sending 50 flows from H4 (10.10.10.115) to H2 (10.10.10.15)

H3 to H1 traffic details:

- SMAC: 00:10:94:00:00:23
- DMAC: 00:10:94:00:00:24
- SIP: 10.10.10.110
- DIP: 10.10.10.10
- SPORT: 1024-1074
- DPORT:50000 - 50050

H4 to H2 traffic details:

- SMAC: 00:10:94:00:00:c2
- DMAC: 00:10:94:00:00:c1
- SIP: 10.10.10.115
- DIP: 10.10.10.15
- SPORT: 1024-1074
- DPORT:50000 - 50050

Operation

Similar to scenario 1a, QFX5110-1 has received H2's MAC address via either QFX5100-1 or QFX5100-2. As discussed earlier, due to aliasing, QFX5110-1 knows that it can reach H2's MAC address via QFX5100-1 and QFX5100-2, even though only one of them sent the Type 2 advertisement.

The QFX5100/QFX5110 can only install VTEP next hops in the PFE and cannot install ESI next hops. This means that, for any given destination in overlay, only one remote VTEP can be selected. To send traffic to the selected VTEP, traffic can be load balanced at the underlay layer through the two spine nodes.

Hence, for a given destination MAC address, there is no load balancing of the flows at the overlay layer; traffic is load balanced only in the underlay.

In this scenario, traffic flows to multiple destination MAC addresses. The PFE still installs only one remote VTEP next hop, however the kernel programs the PFE such a way that it will alternate between the remote VTEPs that are installed, resulting in a certain amount of load balancing at the overlay layer.

In this case, the first 50 flows are sent through QFX5100-2, while the second 50 flows are sent through QFX5100-1.

Verification

The outputs below focus only on the traffic from H4 (10.10.10.115) to H2 (10.10.10.15). Details for the H3-to-H1 traffic flows are shown in Scenario 1a.

SMAC (H4): 00:10:94:00:00:c2
DMAC (H2): 00:10:94:00:00:c1

Check the MAC table on QFX5110-1 for H4 and H2:

```
{master:0}
user@5110-1> show ethernet-switching table | match 00:c2
   bd10                00:10:94:00:00:c2   D        xe-0/0/46.0

{master:0}
user@5110-1> show ethernet-switching table | match 00:c1
   bd10                00:10:94:00:00:c1   DR       esi.1752
00:01:01:01:01:01:01:01:01:01
```

The output shows that QFX5110-1 has learned H4's MAC address through a local interface, and H2's MAC address over ESI 00:01:01:01:01:01:01:01:01:01.

Verify which remote VTEPs that are advertising ESI 00:01:01:01:01:01:01:01:01:01 to QFX5110-1:

```
user@5110-1> show ethernet-switching vxlan-tunnel-end-point esi | find 1752
00:01:01:01:01:01:01:01:01:01 default-switch           1752  131076  esi.1752
2
    RVTEP-IP              RVTEP-IFL      VENH     MASK-ID   FLAGS
    172.23.0.1             vtep.32772     1753      1          2
    172.24.0.1             vtep.32771     1751      0          2
```

The output shows the remote VTEPs and VTEP IFLs related to the ESI.

Now login to the PFE to see the next hop installed for H2's MAC address:

```
FPC0(5110-1 vty)# show l2 manager mac-table

...
mac address          BD     learn  Entry  entry      hal         hardware info
                     Index  vlan   Flags  ifl        ifl         pfe  mask  ifl
  ---------------------------------------------------------------------------------
  00:10:94:00:00:c1  3      0     0x0014 vtep.32772  vtep.32772   0    0x1  vtep.32772
  00:10:94:00:00:c2  3      0     0x0814 xe-0/0/46.0 xe-0/0/46.0  0    0x1  xe-0/0/46.0
...
```

Then verify the remote VTEP's IP address:

```
user@5110-1> show interfaces vtep.32772 | match endpoint
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 172.23.0.1, L2 Routing Instance:
default-switch, L3 Routing Instance: default
```

The output shows that the device has installed the VTEP next hop, and not the ESI next hop. For H2's MAC address, the VTEP installed is 172.23.0.1.

Scenario 2: Inter-VNI traffic using EVPN Type-2 routes

Topology

- Layer 3 gateway at spine layer (QFX10000)
- Layer 2 gateway at leaf layer (QFX5100 / QFX5110)

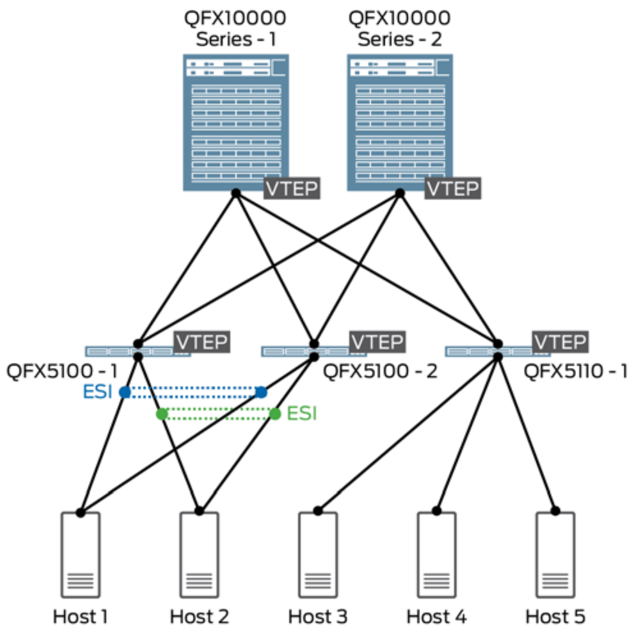

Traffic Flow

- Sending 50 flows from H3 (10.10.10.110) to H5 (172.16.0.5)

Traffic details:

- SMAC: 00:10:94:00:00:23
- DMAC: 00:00:5e:00:01:01 (DMAC of anycast gateway)
- SIP: 10.10.10.110
- DIP: 172.16.0.5
- SPORT: 1024-1074
- DPORT:50000 - 50050

Operation

In this scenario, traffic flows between subnets, therefore it needs to be routed at the spine layer. The packet's DMAC is that of the default gateway, which is on the spine devices. Since anycast gateway (GW) functionality is used, the GW MAC address is advertised with an associated ESI.

QFX5110-1 receives the anycast GW MAC address via either QFX10000-1 or QFX10000-2. Due to aliasing, QFX5110-1 knows that it can reach the GW MAC address via QFX10000-1 and QFX10000-2, even though only one of them sent the Type 2 advertisement.

As discussed in Scenario 1, QFX5110-1 can only install VTEP next hops in the PFE, it cannot install ESI next hops. For inter-VNI traffic, the DMAC is always set to the anycast GW MAC, hence it will always select only one GW.

Once traffic reaches one of the QFX10000 devices, it performs the routing (from VNI 10 to VNI 30 in this case). The destination in VNI 30, i.e. H5, is reachable through both QFX5100-1 and QFX5100-2.

This is where load-balancing behavior differs between QFX10000 Series switches and QFX5100/QFX5110 switches: the QFX10000 will actually install the ESI next hop in the PFE. Therefore, it can load balance flows to a particular destination MAC address across the remote VTEP on QFX5100-1 *and* the remote VTEP on QFX5100-2.

Verification

Check the MAC table on QFX5110-1 for the anycast GW MAC address:

```
{master:0}
user@5110-1> show ethernet-switching table | match 00:00:5e:00:01:01
   bd10                00:00:5e:00:01:01   DR       esi.1744
05:00:00:00:00:00:00:00:0a:00
```

The output shows that QFX5110-1 has learned the anycast GW MAC address via ESI 05:00:00:00:00:00:00:00:0a:00.

Verify which remote VTEPs that are advertising ESI 05:00:00:00:00:00:00:00:0a:00 to QFX5110-1:

```
{master:0}
user@5110-1> show ethernet-switching vxlan-tunnel-end-point esi
ESI                           RTT                       VLNBH INH     ESI-IFL   LOC-IFL
#RVTEPs
05:00:00:00:00:00:00:00:0a:00 default-switch            1744  131081  esi.1744
2
    RVTEP-IP             RVTEP-IFL      VENH     MASK-ID   FLAGS
    172.22.0.1            vtep.32770     1750     1         2
    172.21.0.1            vtep.32769     1742     0         2
```

The output shows the remote VTEPs and VTEP IFLs related to the ESI.

Now login to the PFE to see the next hop installed for the GW MAC address:

```
FPC0(5110-1 vty)# show l2 manager mac-table

...
mac address        BD     learn  Entry  entry      hal        hardware info
                   Index  vlan   Flags  ifl        ifl        pfe  mask  ifl
```

```
  --------------------------------------------------------------------------------
  00:00:5e:00:01:01  3     0     0x0014 vtep.32770 vtep.32770  0   0x1  vtep.32770
...
```

Then verify the remote VTEP's IP address:

```
{master:0}
user@5110-1> show interfaces vtep.32770 | match endpoint
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 172.22.0.1, L2 Routing Instance:
default-switch, L3 Routing Instance: default
```

The output shows that the device has installed the VTEP next hop, and not the ESI next hop. In this scenario, the VTEP installed is 172.22.0.1, which is on QFX10000-2.

Once the traffic reaches QFX10000-2, a lookup is performed to learn the MAC address of H5 (00:10:94:00:00:C0):

```
{master:0}
user@10000-2> show ethernet-switching table | match :c0
   bd30                00:10:94:00:00:c0   DR       esi.1825
00:01:01:01:01:01:01:01:01:01
```

The output shows that QFX10002-2 has learned H5's MAC address via ESI 00:01:01:01:01:01:01:01:01:01.

Verify which remote VTEPs that are advertising ESI 00:01:01:01:01:01:01:01:01:01 to QFX10002-2:

```
{master:0}
user@10000-2> show ethernet-switching vxlan-tunnel-end-point esi
ESI                           RTT                      VLNBH INH     ESI-IFL   LOC-IFL
#RVTEPs
00:01:01:01:01:01:01:01:01:01 default-switch           1825  2097155 esi.1825
2
    RVTEP-IP             RVTEP-IFL      VENH     MASK-ID   FLAGS
    172.23.0.1            vtep.32771     1826     1         2
    172.24.0.1            vtep.32769     1824     0         2
```

The output shows the remote VTEPs and VTEP IFLs related to the ESI.

Now login to the PFE to see the next hop installed for H5's MAC address:

```
FPC0(10000-2 vty)# show l2 manager mac-table

...
mac address          BD     learn  Entry  entry      hal        hardware info
                     Index  vlan   Flags  ifl        ifl        pfe  mask  ifl
  --------------------------------------------------------------------------------

  00:10:94:00:00:c0  11     0     0x0014 esi.1825   esi.1825    0   0x1  esi.1825
...
```

Since the next hop is an ESI, traffic will be load balanced across both QFX5100-1 and QFX5100-2.

Scenario 3: Inter-VNI traffic using EVPN Type-5 routes

Topology

- QFX10000s act as collapsed Layer 2 / Layer 3 gateways
- Reachability is advertised using Type 5 messages

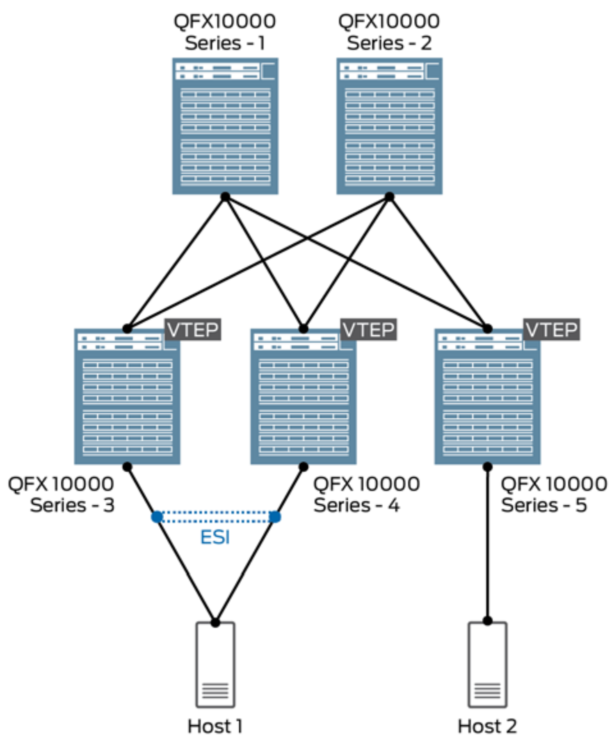

Traffic Flow

- Sending 50 flows from H2 (192.168.210.10) to H1 (192.168.200.10)

Operation

QFX10000-5 receives the 192.168.200.0/24 prefix in a Type 5 route from both QFX10000-3 and QFX10000-4. QFX10000-5 sends all traffic flows to only one destination VTEP. In order to load balance traffic across both remote VTEPs, the `multipath` parameter must be added to the VRF's routing options, as follows:

```
set routing-instances VRF-1 routing-options multipath
```

Verification

Before enabling multipath on QFX10000-5:

```
{master:0}
```



```
user@10000-5> show evpn ip-prefix-database prefix 192.168.200.0/24
L3 context: VRF-1

EVPN->IPv4 Imported Prefixes
Prefix                                             Etag
192.168.200.0/24                                      0
  Route distinguisher    VNI/Label  Router MAC         Nexthop/Overlay GW/ESI
  172.25.0.10:10         5555       ec:3e:f7:87:dc:5a  172.25.0.1
  172.28.0.10:10         5555       80:ac:ac:2e:75:c8  172.28.0.1

{master:0}
user@10000-5> show route forwarding-table | match 192.168.200.0
192.168.200.0/24      user     0                    comp     1831     2
```

The output shows a single composite next hop type for the prefix in the Packet Forwarding Engine, indicating that the device is selecting only one of the two available next hops.

After enabling multipath on QFX10000-5:

```
{master:0}
root@qfx10k-5> show route forwarding-table | find 192.168.200.0
192.168.200.0/24      user     0                    ulst  2097154     2
                                                 comp     1831     2
                                                 comp     1832     2
```

With multipath enabled, the output now shows the unilist next hop type for the prefix with two next hops, indicating that ECMP is taking place.

References

Juniper Networks EVPN Implementation for Next-Generation Data Center Architectures

Configuring EVPN Type 5 for QFX10000 Series Switches

Understanding EVPN Pure Route Type-5 on QFX Series Switches

Configuring EVPN-VXLAN In a Collapsed IP Fabric Topology Within a Data Center

Configuring EVPN for IRB Virtual Gateway Support in EVPN-VXLAN Deployments