



## VXLAN Flood and Learn

« VXLAN Underlay Infrastructure Routing | VXLAN BGP EVPN »

Last updated: January 19, 2017

### VXLAN Flood and Learn

#### Note:

Either configure the section [VXLAN Underlay Infrastructure Routing](#) or load its [solution configs](#) before configuring this task.

#### Task

- Configure your first 5Ks link to your first Server as an edge port and access port in VLAN 172.
- Configure your second 5Ks link to your second Server as an edge port and access port in VLAN 172.
- The servers should use IP addresses 172.16.0.1X/24, where X is the Server number.
- Configure VXLAN Flood and Learn on your 5Ks as follows:
  - VLAN 172 should map to VXLAN Network Identifier (VNI) 1721600.
  - Source the VXLAN Tunnel Endpoint (VTEP) from the 5Ks' Loopback0 interfaces.
  - Use the multicast group 224.172.16.0 for the replication of Broadcast, Unknown, and Multicast (BUM) traffic.
- Once complete, your servers should have layer 2 connectivity to each other over the routed fabric from the Leafs to the Spines, and when you check the MAC table of the Leaf switches, the address of the remote server should be reachable via the Network Virtualization Edge (NVE) tunnel interface.
- Ensure that connectivity is maintained if any single link from the Leafs to the Spines goes down, or if either of the Spines goes down.

#### Configuration [Click to collapse](#)

#### Note:

Nexus 5600 is required to run in store-and-forward switching mode in order to support VXLAN encapsulation. Saving the config and reloading the switch is required in order to activate this configuration change.

```
N5K1# config t
Enter configuration commands, one per line. End with CNTL/Z.
N5K1(config)# hardware ethernet store-and-fwd-switching
Enabling store-and-forward switching. Please copy the configuration and reload the switch
N5K1(config)# end
N5K1# copy run start
[#####] 100%
Copy complete, now saving to disk (please wait)...
N5K1# reload
WARNING: This command will reboot the system
Do you want to continue? (y/n) [n]: y
```

```
N5K1:
feature nv overlay
feature vn-segment-vlan-based
!
vlan 172
  vn-segment 1721600
!
interface Ethernet1/1
  switchport access vlan 172
  spanning-tree port type edge
!
interface nve1
  no shutdown
  source-interface loopback0
  member vni 1721600 mcast-group 224.172.16.0

N5K2:
feature nv overlay
feature vn-segment-vlan-based
!
vlan 172
  vn-segment 1721600
!
interface Ethernet1/2
  switchport access vlan 172
  spanning-tree port type edge
!
interface nve1
  no shutdown
  source-interface loopback0
  member vni 1721600 mcast-group 224.172.16.0
```

#### Verification

VXLAN Flood and Learn uses Multicast replication for VXLAN Tunnel Endpoints (VTEPs) to discover each other, and in order to populate the MAC address to VTEP mappings between all Leaf switches in the VXLAN network.

This is similar to the behavior in a Classical Ethernet network, which uses layer 2 broadcasts to forward frames such as ARP, however in VXLAN these layer 2 broadcasts are first encapsulated in layer 3 multicast. This means that the first verification for VXLAN Flood and Learn is to ensure that the VTEPs can join the same multicast tree.

```
N5K1# show ip mroute 224.172.16.0
IP Multicast Routing Table for VRF "default"
(*, 224.172.16.0/32), bidir, uptime: 02:55:47, nve ip pim
```

<https://t.me/learningnets>

```

Incoming interface: Ethernet1/5, RPF nbr: 10.0.0.71
Outgoing interface list: (count: 2)
  Ethernet1/5, uptime: 02:55:47, pim, (RPF)
  nve1, uptime: 02:55:47, nve

N5K2# show ip mroute 224.172.16.0
IP Multicast Routing Table for VRF "default"

(*, 224.172.16.0/32), bidir, uptime: 02:54:18, nve ip pim
Incoming interface: Ethernet1/5, RPF nbr: 10.0.0.71
Outgoing interface list: (count: 2)
  Ethernet1/5, uptime: 02:54:18, pim, (RPF)
  nve1, uptime: 02:54:18, nve

N7K1# show ip mroute 224.172.16.0
IP Multicast Routing Table for VRF "default"

(*, 224.172.16.0/32), bidir, uptime: 02:55:47, pim ip
Incoming interface: loopback1, RPF nbr: 10.0.0.255
Outgoing interface list: (count: 3)
  Ethernet1/4, uptime: 02:54:18, pim
  loopback1, uptime: 02:55:47, pim, (RPF)
  Ethernet1/3, uptime: 02:55:47, pim

N7K2# show ip mroute 224.172.16.0
IP Multicast Routing Table for VRF "default"

(*, 224.0.0.0/4), bidir, uptime: 07:27:40, pim ip
Incoming interface: loopback1, RPF nbr: 10.0.0.255
Outgoing interface list: (count: 1)
  loopback1, uptime: 07:27:40, pim, (RPF)

```

Note that in the above output the secondary RP, N7K2, will not build the (\*, 224.172.16.0) tree for VNI 1721600 unless the primary RP, N7K1, fails.

Once BUM traffic is flooded (e.g. ARPs from Server 1 to Server 2), the Leaf switches should discover each other as VTEPs. In the below output the DP field means the peer was learned in the Data Plane (Flood and Learn) as opposed to the Control Plane (BGP L2VPN EVPN).

```

N5K1# show nve peers
Interface Peer-IP      State LearnType Uptime  Router-Mac
-----
nve1      10.0.0.52      Up      DP      00:00:41 n/a

N5K1# show nve peers detail
Details of nve Peers:
-----
Peer-IP: 10.0.0.52
NVE Interface      : nve1
Peer State         : Up
Learn Type         : DP
Peer Uptime        : 00:01:04
Router-Mac         : n/a
Peer First VNI     : 1721600
Time since Create  : 00:01:04
Configured VNIs   : 1721600
Provision State    : add-complete
Route-Update      : Yes
Peer Flags        : None
Learn Src         : UDP
Learnt CP VNIs   : --
Peer-ifindex-resp : Yes
-----

```

Once BUM traffic is replicated between the VTEPs over the multicast underlay network, the CAM tables of the Leaf switches should be populated with the MAC address to VTEP mappings, as seen below.

```

N5K1# show mac address-table vlan 172
Legend:
  * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
  age - seconds since last seen,+ - primary entry using vPC Peer-Link
VLAN  MAC Address      Type      age      Secure NTFY  Ports/SWID.SSID.LID
-----
* 172  001b.2188.8075     dynamic   210      F   F   nve1/10.0.0.52
* 172  001b.218d.3d98     dynamic   180      F   F   Eth1/1

N5K2# show mac address-table vlan 172
Legend:
  * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
  age - seconds since last seen,+ - primary entry using vPC Peer-Link
VLAN  MAC Address      Type      age      Secure NTFY  Ports/SWID.SSID.LID
-----
* 172  001b.2188.8075     dynamic   210      F   F   Eth1/2
* 172  001b.218d.3d98     dynamic   180      F   F   nve1/10.0.0.51

```

Layer 2 traffic between the above two MAC addresses follows the IPv4 unicast routed path between VTEP addresses 10.0.0.51 and 10.0.0.52, i.e. the Loopback0 interfaces of the Leaf switches. Since all Leaf to Spine links are the same bandwidth (10GigE), they are the same OSPF cost, and ECMP load distribution from the Leafs to the Spines is achieved.

This can be verified by using the iPerf app on the servers, as seen below.

```

Server 1
C:\Users\Administrator\Desktop\iperf-3.1.3-win64>ipconfig | find "IPv4"
IPv4 Address. . . . . : 192.168.0.11
Autoconfiguration IPv4 Address. . : 169.254.218.42
IPv4 Address. . . . . : 172.16.0.11

C:\Users\Administrator\Desktop\iperf-3.1.3-win64>iperf3.exe -s
-----https://t.me/learningnets
Server listening on 5201

```

## Server 2

```
C:\Users\Administrator\Desktop\iperf-3.1.3-win64>iperf3.exe -c 172.16.0.11 -P 8 -t 10000
Connecting to host 172.16.0.11, port 5201
[ 4] local 172.16.0.12 port 49695 connected to 172.16.0.11 port 5201
[ 6] local 172.16.0.12 port 49696 connected to 172.16.0.11 port 5201
[ 8] local 172.16.0.12 port 49697 connected to 172.16.0.11 port 5201
[10] local 172.16.0.12 port 49698 connected to 172.16.0.11 port 5201
[12] local 172.16.0.12 port 49699 connected to 172.16.0.11 port 5201
[14] local 172.16.0.12 port 49700 connected to 172.16.0.11 port 5201
[16] local 172.16.0.12 port 49701 connected to 172.16.0.11 port 5201
[18] local 172.16.0.12 port 49702 connected to 172.16.0.11 port 5201
[ ID] Interval           Transfer             Bandwidth
[ 4]  0.00-1.01   sec   42.6 MBytes       355 Mbits/sec
[ 6]  0.00-1.01   sec   26.2 MBytes       218 Mbits/sec
[ 8]  0.00-1.01   sec   26.1 MBytes       217 Mbits/sec
[10]  0.00-1.01   sec    256 KBytes       2.08 Mbits/sec
[12]  0.00-1.01   sec   28.6 MBytes       238 Mbits/sec
[14]  0.00-1.01   sec   42.5 MBytes       354 Mbits/sec
[16]  0.00-1.01   sec    256 KBytes       2.08 Mbits/sec
[18]  0.00-1.01   sec   19.1 MBytes       159 Mbits/sec
[SUM] 0.00-1.01   sec   186 MBytes       1.55 Gbits/sec
-----
<snip>
```

8 TCP flows were generated from the servers, each with a unique source TCP port. This traffic then gets encapsulated into VXLAN with a unique source UDP port, and hashed amongst the links in the fabric from Leaf to Spines. The below output shows a near 1:1 ratio of traffic transmitting from Leaf N5K2 out to the Spines, and from Leaf N5K1 receiving traffic from the Spines, with a total throughput of around 2Gbps.

```
N5K2# show interface e1/5-6 | in rate
30 seconds input rate 13193288 bits/sec, 14357 packets/sec
30 seconds output rate 1415783936 bits/sec, 274832 packets/sec
input rate 10.73 Mbps, 11.62 Kpps; output rate 1.15 Gbps, 223.06 Kpps
30 seconds input rate 14858304 bits/sec, 16173 packets/sec
30 seconds output rate 1250529448 bits/sec, 242755 packets/sec
input rate 12.09 Mbps, 13.10 Kpps; output rate 1.01 Gbps, 196.97 Kpps

N5K1# show interface e1/5-6 | in rate
30 seconds input rate 1435753392 bits/sec, 278709 packets/sec
30 seconds output rate 13014336 bits/sec, 14165 packets/sec
input rate 1.12 Gbps, 218.22 Kpps; output rate 10.45 Mbps, 11.32 Kpps
30 seconds input rate 1250194536 bits/sec, 242691 packets/sec
30 seconds output rate 14865032 bits/sec, 16186 packets/sec
input rate 987.35 Mbps, 191.62 Kpps; output rate 11.82 Mbps, 12.81 Kpps
```



[« VXLAN Underlay Infrastructure Routing | VXLAN BGP EVPN »](#)

[^ back to top](#)