



Virtual eXtensible Local-Area Network

(VXLAN)



Virtual Extensible LAN (VXLAN)

- Traditional layer 2 networks have issues because of three main reasons:
 - Spanning-tree.
 - Limited amount of VLANs.
 - Large MAC address tables.
- **Spanning-tree blocks any redundant links to avoid loops.** Blocking links to create a loop-free topology gets the job done, but it also means **we pay for links we can't use.** We could switch to a layer 3 network, but **some technology requires layer 2 networking.**
- The **VLAN ID is 12-bit**, which means we can create 4094 VLANs (0 and 4095 are reserved). Only 4094 available VLANs **can be an issue for data centers.** For example, **imagine we have a service provider with 500 customers. With 4094 available VLANs, they can only offer 8 VLANs to each customer.**




Virtual Extensible LAN (VXLAN)

- **Because of server virtualization**, the number of addresses in the MAC address tables of our switches has grown exponentially. Before server virtualization, a switch only had to learn one MAC address per switchport.
- **With server virtualization**, we run many virtual machines (VM) or containers on a single physical server. Each VM has a virtual NIC and a virtual MAC address. The switch has to learn **many MAC addresses on a single switchport**.
- A Top of Rack (ToR) switch in a data center could connect to 24 or 48 physical servers. A data center could have many racks, so each switch has to store the MAC addresses of all VMs that communicate with each other.
- **We require much larger MAC address tables compared to networks without server virtualization.**

A1	B1	A2	91	23	71
48	51	A2	E1	A3	71
71	44	A2	F1	A9	71
22	B1	55	E1	A3	77


A0	01	72	E9	A8	77
E2	B1	A2	E1	A3	91
A7	B7	82	E1	65	11
88	51	22	66	A3	33

Virtual Extensible LAN (VXLAN)



Virtual Extensible LAN (VXLAN)

- **An overlay network** is a **virtual** network that runs on top of a **physical** underlay network. A GRE tunnel is a simple example of an overlay network. The GRE tunnel runs on top of a physical underlay network.




- **With VXLAN**, the overlay is a **layer 2 Ethernet network**. The underlay network is a **layer 3 IP network**. Another name for the underlay network is a **transport network**.
- The underlay network is simple; its only job is to get packets from A to B. We don't use any layer 2 here, **only layer 3**.
- Another advantage is that the overlay and underlay network are independent. The overlay network is virtual and requires an underlay network, but whatever changes you make in the overlay network won't affect the underlay network. You can add and remove links in the underlay network, and as long as your routing protocol can reach the destination, your overlay network will remain unchanged.

Virtual Extensible LAN (VXLAN)

- **VXLAN** uses 24-Bit VLAN ID (VXLAN Network ID- VNID) instead of 12-Bit VLAN ID.
- This allows 16 Million ID in one administrative domain.
- No Spanning Tree, **VXLAN** utilize L3 infrastructure.




- It is called **MAC in UDP encapsulation**. Tunneling L2 packet over L3 network.



VTEP

- The **VXLAN tunnel endpoint (VTEP)** is the device that's responsible for encapsulating and de-encapsulating layer 2 traffic. **This device is the connection between the overlay and the underlay network.**
- The VTEP comes in two forms:
 - **Software (host-based)** : Hypervisors like VMWare's ESXi or Microsoft's Hyper-V. These hypervisors use virtual switches, and some of them support VXLAN.


The VXLAN tunnels are between the virtual switches of the hypervisors. The underlay network is unaware of VXLAN.



- **Hardware (gateway)** : A hardware VTEP is a router, switch, or firewall which supports VXLAN. We also call a hardware VTEP a **VXLAN gateway** because it combines a regular VLAN and VXLAN segment into a single layer 2 domain.


Some switches have VXLAN support with **ASICs**, offering better VXLAN performance than a software VTEP.

Virtual Extensible LAN (VXLAN)



- Devices that support VXLANs are called virtual tunnel endpoints (VTEPs)

VXLAN Bridging LAB





INET CCNP ENCOR VXLAN - LISP

I-NET ©
Draft and Copyright

VXLAN Bridging LAB

```
VPCS> set pcname Host1
```

```
Host1> ip 10.1.1.1/24
```

```
Checking for duplicate address...
```

```
PC1 : 10.1.1.1 255.255.255.0
```

```
Host1>
```

```
VPCS> set pcname Host2
```

```
Host2> ip 10.1.1.2/24
```

```
Checking for duplicate address...
```

```
PC1 : 10.1.1.2 255.255.255.0
```

```
Host2>
```



VXLAN Bridging LAB

```
switch(config)# hostname Leaf1
```

```
Leaf1(config)# vlan 10
```

```
Leaf1(config-vlan)# exit
```

```
Leaf1(config)# int e1/1
```

```
Leaf1(config-if)# switchport
```

```
Leaf1(config-if)# switchport mode trunk
```

```
Leaf1(config-if)# switchport trunk allowed vlan 10
```

```
Leaf1(config-if)# no shutdown
```

```
Leaf1(config-if)# int e1/2
```

```
Leaf1(config-if)# switchport
```

```
Leaf1(config-if)# switchport mode access
```

```
Leaf1(config-if)# switchport access vlan 10
```

```
Leaf1(config-if)# no shutdown
```

```
Leaf1(config-if)# copy run sta
```

```
[#####] 100%
```

```
Copy complete, now saving to disk (please wait)...
```

```
Copy complete.
```



VXLAN Bridging LAB

```
Leaf1(config)# show run int Ethernet 1/1-2
```

```
!Command: show running-config interface Ethernet1/1-2
```

```
!Running configuration last done at: Tue Jun 11 07:43:03 2024
```

```
!Time: Tue Jun 11 07:47:07 2024
```

```
version 7.0(3)I7(7) Bios: version
```

```
interface Ethernet1/1
```

```
  switchport mode trunk
```

```
  switchport trunk allowed vlan 10
```

```
interface Ethernet1/2
```

```
  switchport access vlan 10
```



VXLAN Bridging LAB

```
Leaf1(config)# show int Ethernet 1/1-2 switchport
```

```
Name: Ethernet1/1
```

```
Switchport: Enabled
```

```
Switchport Monitor: Not enabled
```

```
Switchport Isolated : Not enabled
```

```
Switchport Block Multicast: Not enabled
```

```
Switchport Block Unicast: Not enabled
```

```
Operational Mode: trunk
```

```
Access Mode VLAN: 1 (default)
```

```
Trunking Native Mode VLAN: 1 (default)
```

```
Trunking VLANs Allowed: 10
```

```
Voice VLAN: none
```

```
Extended Trust State : not trusted [COS = 0]
```

```
Administrative private-vlan primary host-association: none
```

```
Administrative private-vlan secondary host-association: none
```

```
Administrative private-vlan primary mapping: none
```

```
Administrative private-vlan secondary mapping: none
```

```
Administrative private-vlan trunk native VLAN: none
```

```
...
```



INET CCNP ENCOR VXLAN - LISP

VXLAN Bridging LAB

Leaf1(config)# show vlan id 10

VLAN	Name	Status	Ports
----	-----	-----	-----
10	VLAN0010	active	Eth1/1, Eth1/2

VLAN Type	Vlan-mode
-----	-----
10 enet	CE

Remote SPAN VLAN

Disabled

Primary	Secondary	Type	Ports
-----	-----	-----	-----



VXLAN Bridging LAB

```
switch(config)# hostname Leaf2
```

```
Leaf2(config)# vlan 10
```

```
Leaf2(config-vlan)# exit
```

```
Leaf2(config)# int e1/1
```

```
Leaf2(config-if)# switchport
```

```
Leaf2(config-if)# switchport mode trunk
```

```
Leaf2(config-if)# switchport trunk allowed vlan 10
```

```
Leaf2(config-if)# no shutdown
```

```
Leaf2(config-if)# int e1/2
```

```
Leaf2(config-if)# switchport
```

```
Leaf2(config-if)# switchport mode access
```

```
Leaf2(config-if)# switchport access vlan 10
```

```
Leaf2(config-if)# no shutdown
```

```
Leaf2(config-if)# copy running-config startup-config
```

```
[#####] 100%
```

```
Copy complete, now saving to disk (please wait)...
```

```
Copy complete.
```



INET CCNP ENCOR VXLAN - LISP

I-NET ©
Draft and Copyright

VXLAN Bridging LAB

Leaf2(config)# **show run int Ethernet 1/1-2**

!Command: show running-config interface Ethernet1/1-2

!Running configuration last done at: Tue Jun 11 08:00:06 2024

!Time: Tue Jun 11 08:04:25 2024

version 7.0(3)I7(7) Bios: version

interface Ethernet1/1

switchport mode trunk

switchport trunk allowed vlan 10

interface Ethernet1/2

switchport access vlan 10



VXLAN Bridging LAB

```
Leaf2(config)# show int Ethernet 1/1-2 switchport
```

```
Name: Ethernet1/1
```

```
Switchport: Enabled
```

```
Switchport Monitor: Not enabled
```

```
Switchport Isolated : Not enabled
```

```
Switchport Block Multicast: Not enabled
```

```
Switchport Block Unicast: Not enabled
```

```
Operational Mode: trunk
```

```
Access Mode VLAN: 1 (default)
```

```
Trunking Native Mode VLAN: 1 (default)
```

```
Trunking VLANs Allowed: 10
```

```
Voice VLAN: none
```

```
Extended Trust State : not trusted [COS = 0]
```

```
Administrative private-vlan primary host-association: none
```

```
Administrative private-vlan secondary host-association: none
```

```
Administrative private-vlan primary mapping: none
```

```
Administrative private-vlan secondary mapping: none
```

```
Administrative private-vlan trunk native VLAN: none
```

```
Administrative private-vlan trunk encapsulation: dot1q
```

```
Administrative private-vlan trunk normal VLANs: none
```

```
Administrative private-vlan trunk private VLANs: none
```

```
...
```



INET CCNP ENCOR VXLAN - LISP

VXLAN Bridging LAB

Leaf2(config)# **show vlan id 10**

VLAN	Name	Status	Ports
-----	-----	-----	-----
10	VLAN0010	active	Eth1/1, Eth1/2

VLAN Type	Vlan-mode
-----	-----
10 enet	CE

Remote SPAN VLAN

Disabled

Primary	Secondary	Type	Ports
-----	-----	-----	-----



INET CCNP ENCOR VXLAN - LISP

VXLAN Bridging LAB

```
switch# conf t
switch(config)# hostname Spine1
```

```
Spine1(config)# vlan 10
```

```
Spine1(config-vlan)# int ethernet 1/1-2
Spine1(config-if-range)# switchport
Spine1(config-if-range)# switchport mode trunk
Spine1(config-if-range)# switchport trunk allowed vlan 10
Spine1(config-if-range)# no shutdown
```

```
Spine1(config-if-range)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.
```



INET CCNP ENCOR VXLAN - LISP

VXLAN Bridging LAB

Spine1(config)# **show run int e1/1-2**

!Command: show running-config interface Ethernet1/1-2

!Running configuration last done at: Tue Jun 11 08:09:44 2024

!Time: Tue Jun 11 08:11:54 2024

version 7.0(3)I7(7) Bios:version

interface Ethernet1/1

switchport mode trunk

switchport trunk allowed vlan 10

interface Ethernet1/2

switchport mode trunk

switchport trunk allowed vlan 10



INET CCNP ENCOR VXLAN - LISP

VXLAN Bridging LAB

Host1> ping 10.1.1.2

```
84 bytes from 10.1.1.2 icmp_seq=1 ttl=64 time=37.988 ms
84 bytes from 10.1.1.2 icmp_seq=2 ttl=64 time=47.058 ms
84 bytes from 10.1.1.2 icmp_seq=3 ttl=64 time=21.787 ms
84 bytes from 10.1.1.2 icmp_seq=4 ttl=64 time=20.843 ms
84 bytes from 10.1.1.2 icmp_seq=5 ttl=64 time=21.799 ms
```


Host1>

Host2> ping 10.1.1.1

```
84 bytes from 10.1.1.1 icmp_seq=1 ttl=64 time=23.860 ms
84 bytes from 10.1.1.1 icmp_seq=2 ttl=64 time=20.587 ms
84 bytes from 10.1.1.1 icmp_seq=3 ttl=64 time=20.807 ms
84 bytes from 10.1.1.1 icmp_seq=4 ttl=64 time=23.917 ms
84 bytes from 10.1.1.1 icmp_seq=5 ttl=64 time=21.331 ms
```


Host2>

VXLAN Bridging LAB



VXLAN Bridging LAB

VXLAN





VXLAN Bridging LAB

- **Step 1:** Build underlay using OSPF,
- **WHY:** Leafs should be able to reach each other's loopback IP's to form the tunnel.
- **Step 2:** Create NVE interface on Leafs.
- **WHY:** This is the tunnel interface. It's the logical representation of a VTEP.
- **Step 3:** Map VLAN ID to VXLAN ID.
- **WHY:** VLAN tag is removed before encapsulation, so we need another identifier to recognize traffic at the remote VTEP.
- **Step 4:** Configure a method to find the remote VTEP to establish the Tunnel.
- **WHY:** This is a very crucial step. A tunnel between the VTEP's cannot establish if they don't know how to find the remote VTEP. (Static Ingress Replication)



VXLAN Bridging LAB

```
Leaf1# conf t
Leaf1(config)# default interface eth1/1

Leaf1(config)# feature ospf

Leaf1(config)# interface Ethernet 1/1
Leaf1(config-if)# no switchport
Leaf1(config-if)# ip address 11.1.1.0/31
Leaf1(config-if)# ip ospf network point-to-point
Leaf1(config-if)# ip router ospf 1 area 0
Leaf1(config-if)# no shutdown
Leaf1(config-if)# exit

Leaf1(config)# interface loopback 0
Leaf1(config-if)# ip address 100.1.1.1/32
Leaf1(config-if)# ip router ospf 1 area 0
Leaf1(config-if)# exit

Leaf1(config)# router ospf 1
Leaf1(config-router)# router-id 1.1.1.1
Leaf1(config-router)# exit
Leaf1(config)#
```

Step 1:



VXLAN Bridging LAB

```
Leaf1(config)# show run interface e1/1-2
```

```
!Command: show running-config interface Ethernet1/1-2
```

```
!Running configuration last done at: Tue Jun 11 08:51:15 2024
```

```
!Time: Tue Jun 11 08:54:03 2024
```

```
version 7.0(3)I7(7) Bios:version
```

```
interface Ethernet1/1
```

```
no switchport
```

```
ip address 11.1.1.0/31
```

```
ip ospf network point-to-point
```

```
ip router ospf 1 area 0.0.0.0
```

```
no shutdown
```

```
interface Ethernet1/2
```

```
switchport access vlan 10
```

Step 1:



VXLAN Bridging LAB

```
Leaf2# conf t
Leaf2(config)# default interface eth1/1

Leaf2(config)# feature ospf

Leaf2(config)# int e1/1
Leaf2(config-if)# no switchport
Leaf2(config-if)# ip address 12.1.1.0/31
Leaf2(config-if)# ip ospf network point-to-point
Leaf2(config-if)# ip router ospf 1 area 0
Leaf2(config-if)# no shutdown
Leaf2(config-if)# exit

Leaf2(config)# interface loopback 0
Leaf2(config-if)# ip address 100.1.1.2/32
Leaf2(config-if)# ip router ospf 1 area 0
Leaf2(config-if)# exit

Leaf2(config)# router ospf 1
Leaf2(config-router)# router-id 2.2.2.2
Leaf2(config-router)# exit
```

Step 1:



VXLAN Bridging LAB

```
Spine1# conf t
Spine1(config)# default interface Ethernet 1/1-2

Spine1(config)# feature ospf

Spine1(config)# interface Ethernet 1/1
Spine1(config-if)# no switchport
Spine1(config-if)# ip address 11.1.1.1/31
Spine1(config-if)# ip ospf network point-to-point
Spine1(config-if)# ip router ospf 1 area 0
Spine1(config-if)# no shutdown
Spine1(config-if)# exit

Spine1(config)# interface Ethernet 1/2
Spine1(config-if)# no switchport
Spine1(config-if)# ip address 12.1.1.1/31
Spine1(config-if)# ip ospf network point-to-point
Spine1(config-if)# ip router ospf 1 area 0
Spine1(config-if)# no shutdown
Spine1(config-if)# exit

Spine1(config)# router ospf 1
Spine1(config-router)# router-id 3.3.3.3
Spine1(config-router)#
```

Step 1:



VXLAN Bridging LAB

Spine1(config)# **show run interface Ethernet 1/1-2**

!Command: show running-config interface Ethernet1/1-2

!Running configuration last done at: Tue Jun 11 09:05:28 2024

!Time: Tue Jun 11 09:10:04 2024

version 7.0(3)I7(7) Bios:version

interface Ethernet1/1

no switchport

ip address 11.1.1.1/31

ip ospf network point-to-point

ip router ospf 1 area 0.0.0.0

no shutdown

interface Ethernet1/2

no switchport

ip address 12.1.1.1/31

ip ospf network point-to-point

ip router ospf 1 area 0.0.0.0

no shutdown

Step 1:



VXLAN Bridging LAB

Leaf1(config)# **show ip ospf neighbors**

OSPF Process ID 1 VRF default

Total number of neighbors: 1

Neighbor ID	Pri	State	Up Time	Address	Interface
3.3.3.3	1	FULL/ -	00:05:37	11.1.1.1	Eth1/1

Leaf1(config)# **show ip route**

IP Route Table for VRF "default"

'*' denotes best ucast next-hop

...

11.1.1.0/31, ubest/mbest: 1/0, attached

*via 11.1.1.0, Eth1/1, [0/0], 00:21:58, direct

11.1.1.0/32, ubest/mbest: 1/0, attached

*via 11.1.1.0, Eth1/1, [0/0], 00:21:58, local

12.1.1.0/31, ubest/mbest: 1/0

*via 11.1.1.1, Eth1/1, [110/80], 00:06:39, ospf-1, intra

100.1.1.1/32, ubest/mbest: 2/0, attached

*via 100.1.1.1, Lo0, [0/0], 00:21:34, local

*via 100.1.1.1, Lo0, [0/0], 00:21:34, direct

100.1.1.2/32, ubest/mbest: 1/0

*via 11.1.1.1, Eth1/1, [110/81], 00:06:37, ospf-1, intra

Step 1:



VXLAN Bridging LAB

Leaf2(config)# **show ip ospf neighbors**

OSPF Process ID 1 VRF default

Total number of neighbors: 1

Neighbor ID	Pri	State	Up Time	Address	Interface
3.3.3.3	1	FULL/ -	00:09:59	12.1.1.1	Eth1/1

Leaf2(config)# **show ip route**

IP Route Table for VRF "default"

'*' denotes best ucast next-hop

...

11.1.1.0/31, ubest/mbest: 1/0

*via 12.1.1.1, Eth1/1, [110/80], 00:10:54, ospf-1, intra

12.1.1.0/31, ubest/mbest: 1/0, attached

*via 12.1.1.0, Eth1/1, [0/0], 00:19:05, direct

12.1.1.0/32, ubest/mbest: 1/0, attached

*via 12.1.1.0, Eth1/1, [0/0], 00:19:05, local

100.1.1.1/32, ubest/mbest: 1/0

*via 12.1.1.1, Eth1/1, [110/81], 00:10:53, ospf-1, intra

100.1.1.2/32, ubest/mbest: 2/0, attached

*via 100.1.1.2, Lo0, [0/0], 00:18:45, local

*via 100.1.1.2, Lo0, [0/0], 00:18:45, direct

Step 1:



VXLAN Bridging LAB

```
Leaf1(config)# ping 100.1.1.2 source 100.1.1.1
PING 100.1.1.2 (100.1.1.2) from 100.1.1.1: 56 data bytes
64 bytes from 100.1.1.2: icmp_seq=0 ttl=253 time=33.432 ms
64 bytes from 100.1.1.2: icmp_seq=1 ttl=253 time=10.173 ms
64 bytes from 100.1.1.2: icmp_seq=2 ttl=253 time=19.966 ms
64 bytes from 100.1.1.2: icmp_seq=3 ttl=253 time=10.939 ms
64 bytes from 100.1.1.2: icmp_seq=4 ttl=253 time=25.48 ms

--- 100.1.1.2 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 10.173/19.998/33.432 ms
```

Step 1:



VXLAN Bridging LAB

```
Leaf1(config)# feature nv overlay
```

```
Leaf1(config)# interface nve1
```

```
Leaf1(config-if-nve)# no shutdown
```

```
Leaf1(config-if-nve)# source-interface loopback 0
```

```
Leaf1(config-if-nve)# exit
```

```
Leaf2(config)# feature nv overlay
```

```
Leaf2(config)# interface nve1
```

```
Leaf2(config-if-nve)# no shutdown
```

```
Leaf2(config-if-nve)# source-interface loopback 0
```

```
Leaf2(config-if-nve)# exit
```

Step 2:



VXLAN Bridging LAB

```
Leaf1(config)# show interface nve1
nve1 is up
admin state is up, Hardware: NVE
MTU 9216 bytes
Encapsulation VXLAN
Auto-mdix is turned off
RX
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
TX
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
Leaf2(config)# show interface nve1
nve1 is up
admin state is up, Hardware: NVE
MTU 9216 bytes
Encapsulation VXLAN
Auto-mdix is turned off
RX
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
TX
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

Step 2:



VXLAN Bridging LAB

```
Leaf1(config)# show running-config vlan 10
```

```
!Command: show running-config vlan 10
```

```
!Running configuration last done at: Tue Jun 11 09:21:48 2024
```

```
!Time: Tue Jun 11 09:26:29 2024
```

```
version 7.0(3)I7(7) Bios:version
```

```
vlan 10
```

```
Leaf1(config)# feature vn-segment-vlan-based
```

```
Leaf1(config)# vlan 10
```

```
Leaf1(config-vlan)# vn-segment 5000
```

```
Leaf1(config-vlan)# exit
```

```
Leaf2(config)# feature vn-segment-vlan-based
```

```
Leaf2(config)# vlan 10
```

```
Leaf2(config-vlan)# vn-segment 5000
```

```
Leaf2(config-vlan)# exit
```

Warning: Enable double-wide arp-ether tcam carving if igmp snooping/Hsrp over vxlan is enabled.
Ignore if tcam carving is already configured.

Step 3:



VXLAN Bridging LAB

Leaf1(config)# **show running-config vlan 10**

!Command: show running-config vlan 10

!Running configuration last done at: Tue Jun 11 09:28:14 2024

!Time: Tue Jun 11 09:31:35 2024

version 7.0(3)I7(7) Bios:version

vlan 10

vlan 10

vn-segment 5000

Leaf1(config)# **show vxlan**

Vlan	VN-Segment
------	------------

=====	=====
-------	-------

10	5000
-----------	-------------

Step 3:



INET CCNP ENCOR VXLAN - LISP

VXLAN Bridging LAB

Leaf2(config)# **show running-config** vlan 10

!Command: show running-config vlan 10

!Running configuration last done at: Tue Jun 11 09:32:31 2024

!Time: Tue Jun 11 09:33:46 2024

version 7.0(3)I7(7) Bios:version

vlan 10

vlan 10

vn-segment 5000

Leaf2(config)# **show vxlan**

Vlan	VN-Segment
------	------------

====	=====
10	5000

Step 3:



INET CCNP ENCOR VXLAN - LISP

VXLAN Bridging LAB

```
Leaf1(config)# show nve peers  
Leaf1(config)#
```

(Nothing)

```
Leaf1(config)# interface nve1  
Leaf1(config-if-nve)# member vni 5000  
Leaf1(config-if-nve-vni)# ingress-replication protocol static  
Leaf1(config-if-nve-vni-ingr-rep)# peer-ip 100.1.1.2  
Leaf1(config-if-nve-vni-ingr-rep)# exit  
Leaf1(config-if-nve-vni)# exit  
Leaf1(config-if-nve)# exit
```

```
Leaf2(config)# interface nve1  
Leaf2(config-if-nve)# member vni 5000  
Leaf2(config-if-nve-vni)# ingress-replication protocol static  
Leaf2(config-if-nve-vni-ingr-rep)# peer-ip 100.1.1.1  
Leaf2(config-if-nve-vni-ingr-rep)# exit  
Leaf2(config-if-nve-vni)# exit  
Leaf2(config-if-nve)# exit
```

Step 4:



INET CCNP ENCOR VXLAN - LISP

VXLAN Bridging LAB

Leaf1(config)# **show nve peers**

Interface	Peer-IP	State	LearnType	Uptime	Router-Mac
-----	-----	----	-----	-----	-----
Nve1	100.1.1.2	Up	DP	00:03:13	n/a

Leaf1(config)# **show nve vni**

Codes: CP - Control Plane DP - Data Plane
 UC - Unconfigured SA - Suppress ARP
 SU - Suppress Unknown Unicast
 Xconn - Crossconnect
 MS-IR - Multisite Ingress Replication

Interface	VNI	Multicast-group	State	Mode	Type [BD/VRF]	Flags
-----	-----	-----	-----	-----	-----	-----
Nve1	5000	UnicastStatic	Up	DP	L2 [10]	

Step 4:



VXLAN Bridging LAB

```
Leaf1(config)# show l2route mac all
```

Flags -(Rmac):Router MAC (Stt):Static (L):Local (R):Remote (V):vPC link
(Dup):Duplicate (Spl):Split (Rcv):Recv (AD):Auto-Delete (D):Del Pending
(S):Stale (C):Clear, (Ps):Peer Sync (O):Re-Originated (Nho):NH-Override
(Pf):Permanently-Frozen

Topology	Mac Address	Prod	Flags	Seq No	Next-Hops

(Nothing)

```
Host1> ping 10.1.1.2
```

```
84 bytes from 10.1.1.2 icmp_seq=1 ttl=64 time=19.888 ms
84 bytes from 10.1.1.2 icmp_seq=2 ttl=64 time=17.937 ms
84 bytes from 10.1.1.2 icmp_seq=3 ttl=64 time=15.504 ms
84 bytes from 10.1.1.2 icmp_seq=4 ttl=64 time=15.638 ms
84 bytes from 10.1.1.2 icmp_seq=5 ttl=64 time=18.945 ms
```

Step 4:



INET CCNP ENCOR VXLAN - LISP

VXLAN Bridging LAB

Leaf1(config)# **show l2route mac all**

Flags -(Rmac):Router MAC (Stt):Static (L):Local (R):Remote (V):vPC link
(Dup):Duplicate (Spl):Split (Rcv):Recv (AD):Auto-Delete (D):Del Pending
(S):Stale (C):Clear, (Ps):Peer Sync (O):Re-Originated (Nho):NH-Override
(Pf):Permanently-Frozen

Topology	Mac Address	Prod	Flags	Seq No	Next-Hops
-----	-----	-----	-----	-----	-----
10	0050.7966.6890	Local	L,	0	Eth1/2
10	0050.7966.6891	Local	L,	0	(0x47000001) nv

Host1> **show ip**

NAME : Host1[1]
IP/MASK : 10.1.1.1/24
GATEWAY : 0.0.0.0
DNS :
MAC : 00:50:79:66:68:90
LPORT : 20000
MTU : 1500

Step 4:



INET CCNP ENCOR VXLAN - LISP

VXLAN Bridging LAB

Leaf2(config)# **show nve peers**

Interface	Peer-IP	State	LearnType	Uptime	Router-Mac
-----	-----	----	-----	-----	-----
nve1	100.1.1.1	Up	DP	00:14:07	n/a

Leaf2(config)# **show nve vni**

Codes: CP - Control Plane DP - Data Plane
 UC - Unconfigured SA - Suppress ARP
 SU - Suppress Unknown Unicast
 Xconn - Crossconnect
 MS-IR - Multisite Ingress Replication

Interface	VNI	Multicast-group	State	Mode	Type [BD/VRF]	Flags
-----	-----	-----	-----	-----	-----	-----
Nve1	5000	UnicastStatic	Up	DP	L2 [10]	

Step 4:



VXLAN Bridging LAB

Leaf2(config)# **show l2route mac all**

Flags -(Rmac):Router MAC (Stt):Static (L):Local (R):Remote (V):vPC link
(Dup):Duplicate (Spl):Split (Rcv):Recv (AD):Auto-Delete (D):Del Pending
(S):Stale (C):Clear, (Ps):Peer Sync (O):Re-Originated (Nho):NH-Override
(Pf):Permanently-Frozen

Topology	Mac Address	Prod	Flags	Seq No	Next-Hops
-----	-----	-----	-----	-----	-----
10	0050.7966.6890	Local	L,	0	(0x47000001) nv
10	0050.7966.6891	Local	L,	0	Eth1/2

Host2> **show ip**

NAME : Host2[1]
IP/MASK : 10.1.1.2/24
GATEWAY : 0.0.0.0
DNS :
MAC : 00:50:79:66:68:91
LPORT : 20000
MTU : 1500

Step 4:



VXLAN Bridging LAB

Leaf2(config)# **show l2route mac all**

Flags -(Rmac):Router MAC (Stt):Static (L):Local (R):Remote (V):vPC link
(Dup):Duplicate (Spl):Split (Rcv):Recv (AD):Auto-Delete (D):Del Pending
(S):Stale (C):Clear, (Ps):Peer Sync (O):Re-Originated (Nho):NH-Override
(Pf):Permanently-Frozen

Topology	Mac Address	Prod	Flags	Seq No	Next-Hops
10	0050.7966.6890	Local	L,	0	(0x47000001) nv
10	0050.7966.6891	Local	L,	0	Eth1/2

Host2> **show ip**

NAME : Host2[1]
IP/MASK : 10.1.1.2/24
GATEWAY : 0.0.0.0
DNS :
MAC : 00:50:79:66:68:91
LPORT : 20000
MTU : 1500

Note: Remote MAC via
VXLAN tunnel interface

Step 4:



VXLAN Bridging LAB

Leaf1(config)# **show running-config interface nve1**

....

```
interface nve1
  no shutdown
  source-interface loopback0
  member vni 5000
    ingress-replication protocol static
    peer-ip 100.1.1.2
```

Leaf2(config)# **show running-config interface nve1**

...

```
interface nve1
  no shutdown
  source-interface loopback0
  member vni 5000
    ingress-replication protocol static
    peer-ip 100.1.1.1
```

VXLAN Bridging LAB

Host1> ping 10.1.1.2 -c 10000


```
84 bytes from 10.1.1.2 icmp_seq=23 ttl=64 time=21.630 ms
84 bytes from 10.1.1.2 icmp_seq=24 ttl=64 time=13.581 ms
84 bytes from 10.1.1.2 icmp_seq=25 ttl=64 time=14.335 ms
10.1.1.2 icmp_seq=26 timeout
10.1.1.2 icmp_seq=27 timeout
10.1.1.2 icmp_seq=28 timeout
10.1.1.2 icmp_seq=29 timeout
10.1.1.2 icmp_seq=30 timeout
10.1.1.2 icmp_seq=31 timeout
10.1.1.2 icmp_seq=32 timeout
10.1.1.2 icmp_seq=33 timeout
10.1.1.2 icmp_seq=34 timeout
84 bytes from 10.1.1.2 icmp_seq=32 ttl=64 time=20.809 ms
84 bytes from 10.1.1.2 icmp_seq=33 ttl=64 time=15.582 ms
84 bytes from 10.1.1.2 icmp_seq=34 ttl=64 time=29.332 ms
```



```
Leaf1(config)# int nve1
Leaf1(config-if-nve)# shutdown
Leaf1(config-if-nve)# no shutdown
Leaf1(config-if-nve)#
```

Done: Example of Layer 2 VNI for intra-subnet VXLAN packet forwarding.

VXLAN Bridging LAB



11.1.1.0/31 (E1/1)

Lo0: 10.25.1.1

VLAN 10 (E1/2)

ip.src == 10.1.1.1

No.	Time	Source	Destination	Protocol	Length	Info
24	71.12880...	10.1.1.1	10.1.1.2	ICMP	148	Echo (ping) request id=0x6cc8, seq=1/256, ttl=64 (reply in 25)
26	72.14782...	10.1.1.1	10.1.1.2	ICMP	148	Echo (ping) request id=0x6dc8, seq=2/512, ttl=64 (reply in 27)
28	73.16284...	10.1.1.1	10.1.1.2	ICMP	148	Echo (ping) request id=0x6ec8, seq=3/768, ttl=64 (reply in 29)
31	74.17826...	10.1.1.1	10.1.1.2	ICMP	148	Echo (ping) request id=0x6fc8, seq=4/1024, ttl=64 (reply in 32)
33	75.19965...	10.1.1.1	10.1.1.2	ICMP	148	Echo (ping) request id=0x70c8, seq=5/1280, ttl=64 (reply in 34)

Frame 33: 148 bytes on wire (1184 bits), 148 bytes captured (1184 bits) on interface eth0_id 0

- Ethernet II, Src: 50:af:20:00:10:07 (50:af:20:00:10:07), Dst: 50:ef:bc:00:11:07 (50:ef:bc:00:11:07)
- Internet Protocol Version 4, Src: 10.1.1.1, Dst: 10.1.1.2
- User Datagram Protocol, Src Port: 57457, Dst Port: 4789

Source Port: 57457
Destination Port: 4789
Length: 114
[Checksum: [missing]]
[Checksum Status: Not present]
[Stream index: 10]
[Timestamps]

- Virtual eXtensible Local Area Network
- Flags: 0x0800, VXLAN Network ID (VNI)
- 0... .. = GBP Extension: Not defined
- = Don't Learn: False
- ... 1... .. = VXLAN Network ID (VNI): True
- 0... .. = Policy Applied: False
- .000 .000 0.00 .000 = Reserved(R): 0x0000
- Group Policy ID: 0
- VXLAN Network Identifier (VNI): 5000
- Reserved: 0

- Ethernet II, Src: Private_66:68:13 (00:50:79:66:68:13), Dst: Private_66:68:14 (00:50:79:66:68:14)
- Internet Protocol Version 4, Src: 10.1.1.1, Dst: 10.1.1.2
- Internet Control Message Protocol

```

0000  50 ef bc 00 11 07 50 af 20 00 10 07 08 00 45 00  P...P...E-
0010  00 86 80 00 00 00 fe 11 72 61 64 01 01 64 01    ....rad...d
0020  01 02 e0 71 12 b5 00 72 00 00 08 00 00 00 13    ...q...r...
0030  88 00 00 50 79 66 68 14 00 50 79 66 68 13 08 00  ...Pyfh...Pyfh...
  
```


Wireshark_eth0_20260419121841_z7VWMO.pcapng

Packets: 83 · Displayed: 5 (6.0%)

Profile: Default


VXLAN Bridging LAB

VXLAN



What happens if Different VLAN's are mapped on to a VNI?

VXLAN Bridging LAB Ciscov1000



VXLAN Bridging LAB Ciscov1000

```


CSR-8(config)#interface Loopback0
CSR-8(config-if)# ip address 1.1.1.1 255.255.255.255

CSR-8(config-if)#interface GigabitEthernet1
CSR-8(config-if)# mtu 1600
CSR-8(config-if)# ip address 172.16.1.1 255.255.255.252
CSR-8(config-if)#no shutdown
CSR-8(config-if)#exit

CSR-8(config)#ip route 2.2.2.2 255.255.255.255 172.16.1.2
CSR-8(config)#ip route 3.3.3.3 255.255.255.255 172.16.1.2

CSR-8(config)#interface GigabitEthernet4
CSR-8(config-if)# no shutdown
CSR-8(config-if)# no ip address

```




VXLAN Bridging LAB Ciscov1000

```
CSR-8(config)#interface nve 1
CSR-8(config-if)# no ip address
CSR-8(config-if)# no shutdown
CSR-8(config-if)# source-interface Loopback0
CSR-8(config-if)# member vni 10000
CSR-8(config-if-nve-vni)# ingress-replication 2.2.2.2
CSR-8(config-if-nve-vni)# ingress-replication 3.3.3.3
CSR-8(config-if-nve-vni)#exit
CSR-8(config-if)#exit
CSR-8(config)#
```

```
CSR-8(config)#bridge-domain 10
CSR-8(config-bdomain)# member vni 10000
CSR-8(config-bdomain)# member GigabitEthernet4 service
CSR-8(config-bdomain-efp)#exit
CSR-8(config-bdomain)#exit
CSR-8(config)#
```

```
CSR-8(config)#interface GigabitEthernet4
CSR-8(config-if)# service instance 1 ethernet
CSR-8(config-if-srv)# encapsulation untagged
```





INET CCNP ENCOR VXLAN - LISP

VXLAN Bridging LAB Ciscov1000

VPCS> **show ip**

```
NAME           : VPCS[1]
IP/MASK         : 192.168.10.12/24
GATEWAY        : 0.0.0.0
DNS            :
MAC            : 00:50:79:66:68:20
LPORT          : 20000
RHOST:PORT     : 127.0.0.1:30000
MTU            : 1500
```

VPCS> **ping 192.168.10.10**

```
84 bytes from 192.168.10.10 icmp_seq=1 ttl=64 time=10.835 ms
84 bytes from 192.168.10.10 icmp_seq=2 ttl=64 time=5.459 ms
84 bytes from 192.168.10.10 icmp_seq=3 ttl=64 time=5.718 ms
84 bytes from 192.168.10.10 icmp_seq=4 ttl=64 time=5.348 ms
84 bytes from 192.168.10.10 icmp_seq=5 ttl=64 time=5.243 ms
```



INET CCNP ENCOR VXLAN - LISP

VXLAN Bridging LAB Ciscov1000

VPCS> ping 192.168.10.11

84 bytes from 192.168.10.11 icmp_seq=1 ttl=64 time=4.419 ms

84 bytes from 192.168.10.11 icmp_seq=2 ttl=64 time=5.022 ms

...

CSR-33#show bridge-domain 10

Bridge-domain 10 (2 ports in all)

State: UP Mac learning: Enabled

Aging-Timer: 300 second(s)


GigabitEthernet4 service instance 1

vni 10000

AED	MAC address	Policy	Tag	Age	Pseudoport
0	AABB.CC00.0300	forward	dynamic	278	nve1.VNI10000, VxLAN src: 3.3.3.3 dst: 2.2.2.2
0	0050.7966.6819	forward	dynamic	238	nve1.VNI10000, VxLAN src: 3.3.3.3 dst: 2.2.2.2
0	0050.7966.6820	forward	dynamic	238	GigabitEthernet4.EFP1
0	0050.7966.6818	forward	dynamic	231	nve1.VNI10000, VxLAN src: 3.3.3.3 dst: 1.1.1.1
0	AABB.CC00.0110	forward	dynamic	276	GigabitEthernet4.EFP1
0	AABB.CC00.0210	forward	dynamic	277	nve1.VNI10000, VxLAN src: 3.3.3.3 dst: 1.1.1.1


The Simple Breakdown

- **Bridge-Domain (The Local Room):** This is where your local devices sit. It's like a local VLAN that lives inside the router.
- **VNI (The Shipping Label):** This is a unique ID (e.g., 10000) that tells the network, "This packet belongs to the 192.168.10.0 network." It's how the router distinguishes one customer/network from another.
- **NVE Interface (The Tunnel Entrance):** This is the "Portal." It takes the local traffic from the Bridge-Domain, sticks the VNI label on it, and wraps it in a VXLAN header to send it across the internet/underlay.




The "Travel Path" of a Packet

- A packet leaves the PC and hits **Gi4**.
- The router sees it's in **Bridge-Domain 10**.
- The Bridge-Domain sees it's mapped to **VNI 10000**.
- **The NVE 1** interface sees VNI 10000, wraps it in a VXLAN header (UDP 4789), and sends it to the remote Peer IP.
- **Summary:** The Bridge-Domain **collects the traffic**, the VNI **identifies the traffic**, and the NVE **transports the traffic**.



VXLAN Bridging LAB Ciscov1000



Virtual eXtensible Local Area Network (vxlan), 8 byte(s)

Packets: 84 · Displayed: 76 (90.5%)

Profile: Default



Locator-ID Separation Protocol (LISP)


Locator-ID Separation Protocol

- **Internet routing tables have grown exponentially, putting a burden on BGP routers.** Routing on the Internet is meant to be hierarchical, but because of disaggregation, a full Internet routing table **nowadays contains over 1.000.000 prefixes.**
- <https://www.cidr-report.org/as2.0/>
- Disaggregation is the opposite of aggregation (route summarization). We inject more specific routes **when there is an aggregate (summary route).** There are two main reasons why this happens:
 - **Multihoming:** Customers connect to **two different ISPs** and **advertise their provider-independent address space (PI) to both ISPs.**
 - **Traffic engineering:** A common practice for ingress traffic engineering is **to advertise a more specific route.** This works, but it increases the size of the Internet routing table.




LISP Infrastructure Devices

- **MS:** The Map-Server is deployed as a LISP Infrastructure component. It configures the LISP site policy for LISP ETRs that register to it.
- **MR:** The Map-Resolver is deployed as a LISP Infrastructure device.




LISP Control Plane

- The **LISP control plane** is similar to how DNS works:
 - DNS resolves a hostname to an IP address.
 - LISP resolves an EID to an RLOC.**



LISP Configuration



LISP Configuration

R5-MSMR#**sh run**

Building configuration...

!

interface Loopback0

ip address 5.5.5.5 255.255.255.255

!

interface Ethernet0/2

ip address 1.1.1.14 255.255.255.252


!

router ospf 1

router-id 5.5.5.5

network 1.1.1.12 0.0.0.3 area 0

network 5.5.5.5 0.0.0.0 area 0





LISP Configuration

RLOC-SITE-A#show ip route 192.168.2.0

% Network not in table

RLOC-SITE-A#show ip route 192.168.2.102

% Network not in table

RLOC-SITE-B#show ip route 192.168.1.0

% Network not in table

RLOC-SITE-B#show ip route 192.168.1.101

% Network not in table

R5-MSMR(config)#router lisp

R5-MSMR(config-router-lisp)#ipv4 map-resolver

R5-MSMR(config-router-lisp)#ipv4 map-server

R5-MSMR(config-router-lisp)#site R1

R5-MSMR(config-router-lisp-site)#eid-prefix 192.168.1.0/24

R5-MSMR(config-router-lisp-site)#authentication-key INETpassword

R5-MSMR(config-router-lisp-site)#exit

R5-MSMR(config-router-lisp)#site R4

R5-MSMR(config-router-lisp-site)#eid-prefix 192.168.2.0/24

R5-MSMR(config-router-lisp-site)#authentication-key INETpassword



INET CCNP ENCOR VXLAN - LISP

LISP Configuration

R5-MSMR#**show lisp site**

LISP Site Registration Information

Site Name	Last Register	Up	Who Last Registered	Inst ID	EID Prefix
R1	00:00:14	yes	1.1.1.1		192.168.1.0/24
R4	00:00:46	yes	1.1.1.10		192.168.2.0/24

RLOC-SITE-A#**show ip interface brief | exc una**

Interface	IP-Address	OK?	Method	Status	Protocol
Ethernet0/0	1.1.1.1	YES	manual	up	up
Ethernet0/1	192.168.1.1	YES	manual	up	up
LISP0	192.168.1.1	YES	unset	up	up

RLOC-SITE-B#**show ip interface brief | exc una**

Interface	IP-Address	OK?	Method	Status	Protocol
Ethernet0/0	1.1.1.10	YES	manual	up	up
Ethernet0/1	192.168.2.1	YES	manual	up	up
LISP0	192.168.2.1	YES	unset	up	up

LISP Configuration

EID101> ping 192.168.2.102

192.168.2.102 icmp_seq=1 timeout

192.168.2.102 icmp_seq=2 timeout


84 bytes from 192.168.2.102 icmp_seq=3 ttl=60 time=2.678 ms

84 bytes from 192.168.2.102 icmp_seq=4 ttl=60 time=2.460 ms

84 bytes from 192.168.2.102 icmp_seq=5 ttl=60 time=2.953 ms

RLOC-SITE-A#show ip route 192.168.2.0

% Network not in table



LISP Configuration

The image shows a Wireshark packet capture of a LISP Map-Register message. The packet list at the top shows a single packet (No. 10) from 10.14.7.242 to 1.1.1.10, identified as LISP, 106 Map-Register for 192.168.2.0/24.

The packet details pane shows the following structure:

- Frame 10: 106 bytes on wire (848 bits), 106 bytes captured (848 bits) on interface eth0, id 0
- Ethernet II, Src: aa:bb:cc:00:04:00 (aa:bb:cc:00:04:00), Dst: aa:bb:cc:00:03:00 (aa:bb:cc:00:03:00)
- Internet Protocol Version 4, Src: 1.1.1.10, Dst: 5.5.5.5
- User Datagram Protocol, Src Port: 4342, Dst Port: 4342
- Locator/ID Separation Protocol
 - Type: Map-Register (3)
 - 0011 = P bit (Proxy-Map-Reply): Not set
 -0..... = S bit (LISP-SEC capable): Not set
 -0..... = I bit (xTR-ID present): Not set
 -0..... = R bit (Built for an RTR): Not set
 -0000 0000 0000 0000 = Reserved bits: 0x0000
 -0..... = M bit (Want-Map-Notify): Not set
 - Record Count: 1
 - Nonce: 0xc90306f238d302
 - Key ID: 0x0001
 - Authentication Data Length: 20
 - Authentication Data: 7f2929a7f662d9190b5b78ec74553896f8e8fb
 - Mapping Record 1, EID Prefix: 192.168.2.0/24, TTL: 1440, Action: No-Action, Authoritative
 - Record TTL: 1440
 - Locator Count: 1
 - EID Mask Length: 24
 - 000 = Action: No-Action (0)
 -1..... = Authoritative bit: Set
 -0000 0000 = Reserved: 0x000
 -0000 0000 = Reserved: 0x0
 -0000 0000 = Mapping Version: 0
 - EID Prefix AFI: IPv4 (1)
 - EID Prefix: 192.168.2.0
 - Locator Record 1, Local RLOC: 1.1.1.10, Reachable, Priority/Weight: 100/50, Multicast Priority/Weight: 255/0
 - Priority: 100
 - Weight: 50
 - Multicast Priority: 255
 - Multicast Weight: 0
 - Flags: 0x0005
 - AFI: IPv4 (1)
 - Locator: 1.1.1.10

The packet bytes pane at the bottom shows the raw data in hexadecimal and ASCII, with a red box highlighting the authentication data and mapping record fields.

LISP Configuration

The image shows a Wireshark packet capture of an ICMP Echo (ping) request. The packet list shows a request from 192.168.1.101 to 192.168.2.102. The packet details pane shows the following information:

- User Datagram Protocol, Src Port: 1539, Dst Port: 4341**
- Locator/ID Separation Protocol (Data)**
 - Flags: 0x00
 - .1... = N bit (Nonce present): Set
 - .1... = L bit (Locator-Status-Bits field enabled): Set
 - .0... = E bit (Echo-Nonce-Request): Not set
 - ...8... = V bit (Map-Version present): Not set
 - ...0... = I bit (Instance ID present): Not set
 - ...000 = Reserved: 0x0
 - Nonce: 4924018 (0x4b2272)
 - 0000 0000 0000 0000 0000 0000 0001 = Locator-Status-Bits: 0x00000001
- Internet Protocol Version 4, Src: 192.168.1.101, Dst: 192.168.2.102**
 - 0100... = Version: 4
 - ...0101 = Header Length: 20 bytes (5)
 - Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
 - Total Length: 84
 - Identification: 0x54bd (21693)
 - Flags: 0x0000
 - Fragment offset: 0
 - Time to live: 63
 - Protocol: ICMP (1)
 - Header checksum: 0xa1d0 [validation disabled]
 - [Header checksum status: Unverified]
 - Source: 192.168.1.101
 - Destination: 192.168.2.102
- Internet Control Message Protocol**
 - Type: 8 (Echo (ping) request)
 - Code: 0
 - Checksum: 0x62b6 [correct]
 - [Checksum Status: Good]
 - Identifier (BE): 48468 (0xbd54)
 - Identifier (LE): 21693 (0x54bd)
 - Sequence number (BE): 1 (0x0001)
 - Sequence number (LE): 256 (0x0100)
 - [Response frame: 65]
 - Data (56 bytes)

The packet bytes pane shows the raw data of the packet, including the IP header and ICMP data.



INET CCNP ENCOR VXLAN - LISP

I-NET ©
Draft and Copyright

VXLAN_LISP



QUESTIONS?

Prepared by Ayhan KOŞAR

ayhankosar@gmail.com

www.inetworkeskisehir.com

ayhan.k@inetworkeskisehir.com

linkedin.com/in/ayhankosar

