

在Nexus 7K上配置VXLAN泛洪和学习

目录

[简介](#)

[先决条件](#)

[要求](#)

[使用的组件](#)

[配置](#)

[网络图](#)

[配置](#)

[VTEP-1](#)

[VTEP-2](#)

[验证](#)

[输出示例](#)

[VTEP-1](#)

[VTEP-2](#)

[故障排除](#)

简介

本文档介绍Nexus 7000系列交换机上虚拟可扩展局域网(VXLAN)泛洪和学习的配置。

先决条件

要求

Cisco 建议您了解以下主题：

- 组播路由概念，例如交汇点 (RP) 和平台无关组播 (PIM)。
- VXLAN概念

注意：本文档假设，在配置 VXLAN 前已建立 IP 路由和组播路由。

使用的组件

本文档中的信息基于以下软件和硬件版本：

- N77-C7710
- N77-F348XP-23
- N77-F324FQ-25

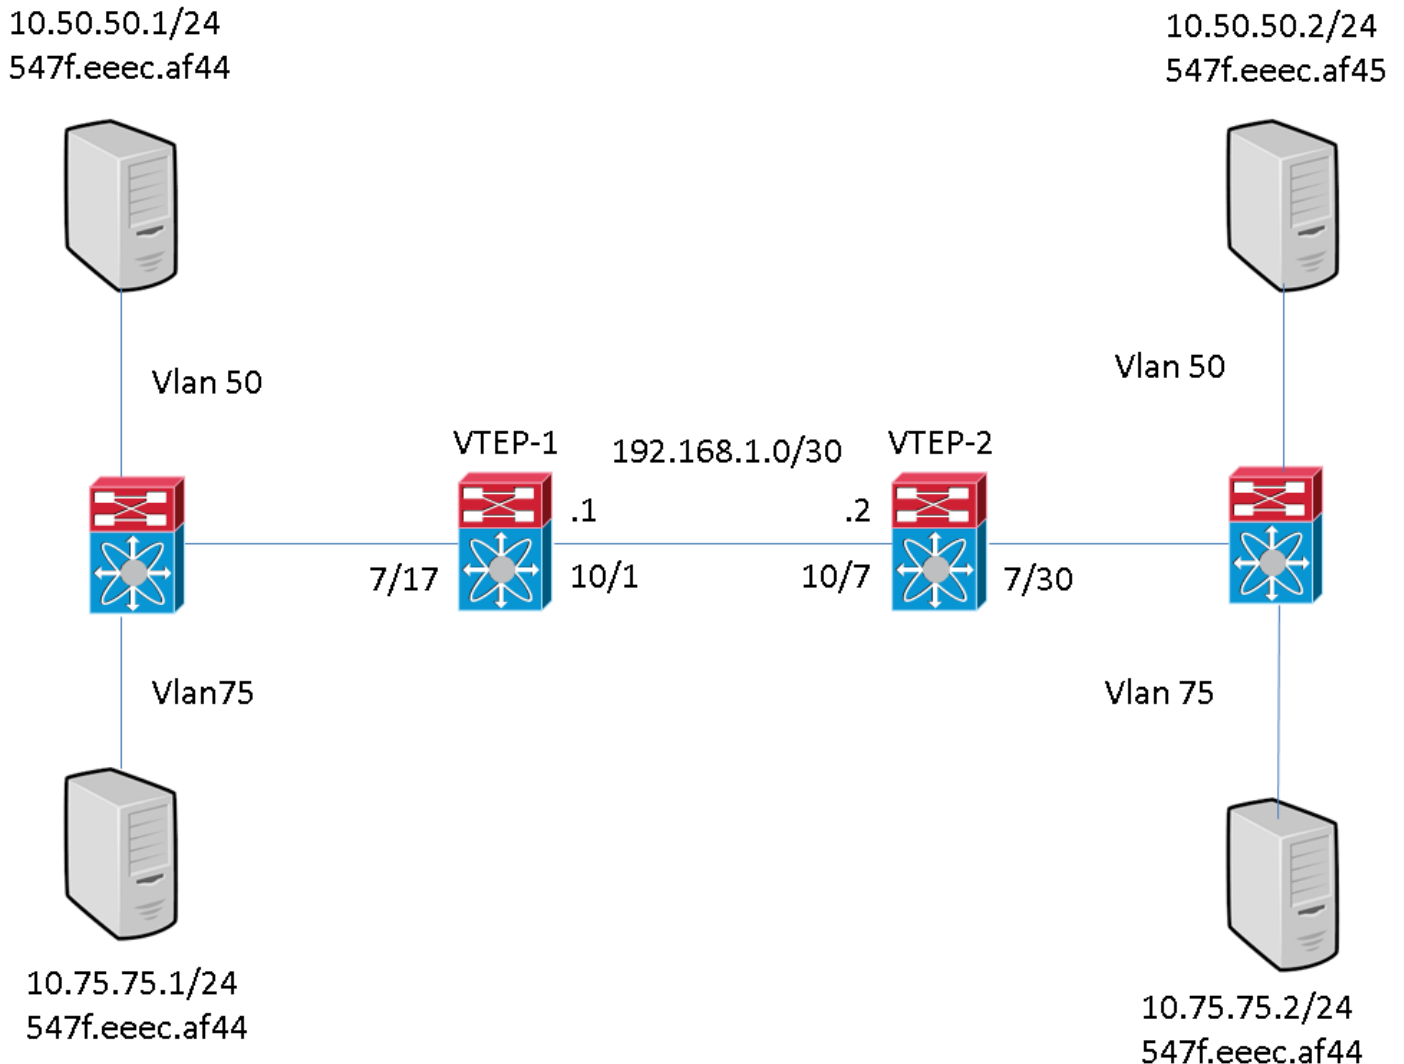
注意：N77K运行软件版本7.2(0)D1(1)。

本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原

始（默认）配置。如果您使用的是真实网络，请确保您已经了解所有命令的潜在影响。

配置

网络图



配置

这些配置是 VXLAN 配置部分的特定配置。这些配置假设，通过您选择的路由协议，可以完全访问拓扑中的所有第 3 层接口。本示例中使用静态路由。它还假设已在这些相同的 L3 接口上建立组播路由

VTEP-1

```
feature pim
system bridge-domain 50,75
feature nv overlay
feature interface-vlan feature vni vni 5000
vni 7500 ip route 10.10.10.2/32 Ethernet10/1 192.168.1.2 ip pim rp-address 192.168.1.1 group-
list 224.0.0.0/4 bridge-domain 50
bridge-domain 75 encapsulation profile vni VSI_50_TO_5000 dot1q 50 vni 5000
encapsulation profile vni VSI_75_TO_7500
```

```

dot1q 75 vni 7500 bridge-domain 50 member vni 5000
bridge-domain 75
member vni 7500 interface nve1 no shutdown source-interface loopback10 member vni 5000 mcast-
group 225.1.1.1
member vni 7500 mcast-group 227.1.1.1

```

```

interface Bdi50
no shutdown
ip address 10.50.50.50/24

```

```

interface Bdi75
no shutdown
ip address 10.75.75.75/24 interface Ethernet7/17
no switchport no shutdown service instance 1 vni no shutdown encapsulation profile
VSI_50_TO_5000 default
service instance 2 vni
no shutdown
encapsulation profile VSI_75_TO_7500 default interface Ethernet10/1
no switchport ip address 192.168.1.1/30 ip pim sparse-mode no shutdown interface loopback10 ip
address 10.10.10.1/32 ip pim sparse-mode

```

请注意，VTEP (Vxlan隧道终端) 上的内部接口配置为第3层端口 (无交换机端口)。但是，没有为其分配IP。另外，必须注意，在VTEP上定义的BD值不必与用于向此设备发送流量的vlan ID匹配。但是，封装配置文件中定义的dot1q到VNI (Vxlan网络标识符) 映射 (在内部接口上的服务实例下调用) 应与Vlan ID匹配。

VTEP-2

```

feature pim
system bridge-domain 50,75
feature nv overlay
feature interface-vlan feature vni vni 5000
vni 7500 ip route 10.10.10.1/32 Ethernet10/7 192.168.1.1 ip pim rp-address 192.168.1.1 group-
list 224.0.0.0/4 bridge-domain 50
bridge-domain 75 encapsulation profile vni VSI_50_TO_5000 dot1q 50 vni 5000
encapsulation profile vni VSI_75_TO_7500
dot1q 75 vni 7500 bridge-domain 50 member vni 5000
bridge-domain 75
member vni 7500 interface nve1 no shutdown source-interface loopback10 member vni 5000 mcast-
group 225.1.1.1
member vni 7500 mcast-group 227.1.1.1

```

```

interface Bdi50
no shutdown
ip address 10.50.50.51/24

```

```

interface Bdi75
no shutdown
ip address 10.75.75.76/24 interface Ethernet7/30
no switchport no shutdown service instance 1 vni no shutdown encapsulation profile
VSI_50_TO_5000 default
service instance 2 vni
no shutdown
encapsulation profile VSI_75_TO_7500 default interface Ethernet10/7
no switchport ip address 192.168.1.2/30 ip pim sparse-mode no shutdown interface loopback10 ip
address 10.10.10.2/32 ip pim sparse-mode

```

请注意，VTEP上的内部接口配置为第3层端口 (无交换机端口)。但是，没有为其分配IP。另外，必须注意，在VTEP上定义的BD值不必与用于向此设备发送流量的Vlan ID匹配。但是，封装配置文件中定义的dot1q到VNI映射 (在内部接口上的服务实例下调用) 应与Vlan ID匹配。

验证

使用本部分可确认配置能否正常运行。

输出示例

这些输出都处于稳定状态。VTEP 对等体已发现彼此，且二者间已在封装和解封两个方向传输流量。

VTEP-1

```
VTEP-1# show nve vni
```

```
Codes: CP - Control Plane      DP - Data Plane
       UC - Unconfigured       SA - Suppress ARP
```

Interface	VNI	Multicast-group	State	Mode	Type [BD/VRF]	Flags
nve1	5000	225.1.1.1	Up	DP	L2 [50]	
nve1	7500	227.1.1.1	Up	DP	L2 [75]	

```
VTEP-1# show running-config interface nve 1
```

```
interface nve1
  no shutdown
  source-interface loopback10
  member vni 5000 mcast-group 225.1.1.1
  member vni 7500 mcast-group 227.1.1.1
```

```
VTEP-1# show service instance vni detail
```

```
VSI: VSI-Ethernet7/17.1
If-index: 0x35310001
Admin Status: Up
Oper Status: Up
Auto-configuration Mode: No
encapsulation profile vni VSI_50_TO_5000
  dot1q 50 vni 5000
Dot1q  VNI    BD
-----
50      5000   50
```

```
VSI: VSI-Ethernet7/17.2
If-index: 0x35310002
Admin Status: Up
Oper Status: Up
Auto-configuration Mode: No
encapsulation profile vni TEST
  dot1q 100 vni 7500
Dot1q  VNI    BD
-----
100     7500   75
```

```
VTEP-1# show bridge-domain
```

```
Bridge-domain 50 (2 ports in all)
Name:: Bridge-Domain50
Administrative State: UP          Operational State: UP
VSI-Eth7/17.1
```

vni5000
nve1

Bridge-domain 75 (2 ports in all)

Name:: Bridge-Domain75

Administrative State: UP

Operational State: UP

VSI-Eth7/17.2

vni7500

nve1

VTEP-1# show mac address-table dynamic

Note: MAC table entries displayed are getting read from software.

Use the 'hardware-age' keyword to get information related to 'Age'

Legend:

* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC

age - seconds since last seen, + - primary entry using vPC Peer-Link, E -

EVPN entry

(T) - True, (F) - False, ~~~ - use 'hardware-age' keyword to retrieve

age info

VLAN/BD	MAC Address	Type	age	Secure NTFY	Ports/SWID.	SSID.LID	-----+-----
nve1/10.10.10.2	* 50 547f.eeec.af44	dynamic	~~~	F F	VSI-Eth7/17.1	* 50 547f.eeec.af45	dynamic ~~~ F F
nve1/10.10.10.2	* 75 547f.eeec.af44	dynamic	~~~	F F	VSI-Eth7/17.2	* 75 547f.eeec.af45	dynamic ~~~ F F

VTEP-1# show ip mroute detail IP Multicast Routing Table for VRF "default" Total number of routes: 7 Total number of (*,G) routes: 2 Total number of (S,G) routes: 4 Total number of (*,G-prefix) routes: 1 (*, 225.1.1.1/32), uptime: 19:51:28, nve(1) ip(0) pim(1) Data Created: No VXLAN Flags VXLAN Encap Stats: 0/0 [Packets/Bytes], 0.000 bps Incoming interface: Ethernet10/1, RPF nbr: 1.1.1.1 Outgoing interface list: (count: 2) Ethernet10/1, uptime: 19:51:09, pim, (RPF) nve1, uptime: 19:51:28, nve (10.10.10.1/32, 225.1.1.1/32), uptime: 19:51:28, nve(0) mrib(0) ip(0) pim(1) Data Created: No Received Register stop VXLAN Flags VXLAN Encap Stats: 19/2274 [Packets/Bytes], 0.000 bps Incoming interface: loopback10, RPF nbr: 10.10.10.1, internal Outgoing interface list: (count: 1) Ethernet10/1, uptime: 19:51:09, pim (10.10.10.2/32, 225.1.1.1/32), uptime: 18:10:06, pim(1) mrib(1) ip(0) Data Created: Yes VXLAN Flags VXLAN Decap Stats: 9/846 [Packets/Bytes], 0.000 bps Incoming interface: Ethernet10/1, RPF nbr: 1.1.1.2, internal Outgoing interface list: (count: 2) Ethernet10/1, uptime: 01:00:32, pim, (RPF) nve1, uptime: 18:10:06, mrib (*, 227.1.1.1/32), uptime: 12:52:13, nve(1) ip(0) pim(1) Data Created: No VXLAN Flags VXLAN Encap Stats: 0/0 [Packets/Bytes], 0.000 bps Incoming interface: Ethernet10/1, RPF nbr: 1.1.1.1 Outgoing interface list: (count: 2) Ethernet10/1, uptime: 12:51:52, pim, (RPF) nve1, uptime: 12:52:13, nve (10.10.10.1/32, 227.1.1.1/32), uptime: 12:52:13, nve(0) mrib(0) ip(0) pim(1) Data Created: No Received Register stop VXLAN Flags VXLAN Encap Stats: 300/39850 [Packets/Bytes], 0.000 bps Incoming interface: loopback10, RPF nbr: 10.10.10.1, internal Outgoing interface list: (count: 1) Ethernet10/1, uptime: 12:51:52, pim (10.10.10.2/32, 227.1.1.1/32), uptime: 12:51:34, pim(1) mrib(1) ip(0) Data Created: Yes VXLAN Flags VXLAN Decap Stats: 22/1928 [Packets/Bytes], 0.000 bps Incoming interface: Ethernet10/1, RPF nbr: 1.1.1.2, internal Outgoing interface list: (count: 2) Ethernet10/1, uptime: 00:52:14, pim, (RPF) nve1, uptime: 12:51:34, mrib (*, 232.0.0.0/8), uptime: 20:56:33, pim(0) ip(0) Data Created: No Stats: 0/0 [Packets/Bytes], 0.000 bps Incoming interface: Null, RPF nbr: 0.0.0.0 Outgoing interface list: (count: 0) VTEP-1# show ip arp Flags: * - Adjacencies learnt on non-active FHRP router + - Adjacencies synced via CFSofE # - Adjacencies Throttled for Glean D - Static Adjacencies attached to down interface IP ARP Table for context default Total number of entries: 4 Address Age MAC Address Interface 10.50.50.1 00:11:32 547f.eeec.af44 Bdi50 10.50.50.2 00:11:14 547f.eeec.af44 Bdi50 10.75.75.1 00:10:45 547f.eeec.af44 Bdi75 10.75.75.2 00:15:04 547f.eeec.af45 Bdi75 192.168.1.2 00:05:39 547f.eeec.af43 Ethernet10/1 VTEP-1# show ip route IP Route Table for VRF "default" '*' denotes best ucast next-hop '**' denotes best mcast next-hop '[x/y]' denotes [preference/metric] '%<string>' in via output denotes VRF <string> 192.168.1.0/30, ubest/mbest: 1/0, attached *via 1.1.1.1, Eth10/1, [0/0], 20:25:13, direct 192.168.1.1/32, ubest/mbest: 1/0, attached *via 1.1.1.1, Eth10/1, [0/0], 20:25:13, local 10.10.10.1/32, ubest/mbest: 2/0, attached *via 10.10.10.1, Lo10, [0/0], 20:25:45, local *via 10.10.10.1, Lo10, [0/0], 20:25:45, direct 10.10.10.2/32, ubest/mbest: 1/0 *via 1.1.1.2, Eth10/1, [1/0], 20:23:42, static 50.50.50.0/24, ubest/mbest: 1/0, attached *via 50.50.50.50, Bdi50, [0/0], 01:18:47, direct 50.50.50.50/32, ubest/mbest: 1/0, attached *via 50.50.50.50, Bdi50, [0/0], 01:18:47, local 75.75.75.0/24, ubest/mbest: 1/0, attached *via 75.75.75.75, Bdi75, [0/0],

01:10:05, direct 75.75.75.75/32, ubest/mbest: 1/0, attached *via 75.75.75.75, Bdi75, [0/0],
01:10:05, local

注意：所有这些输出都通过拓扑中所有主机之间的全网状流量收集。

VTEP-2

VTEP-2# show nve vni

Codes: CP - Control Plane DP - Data Plane
 UC - Unconfigured SA - Suppress ARP

Interface	VNI	Multicast-group	State	Mode	Type	[BD/VRF]	Flags
nve1	5000	225.1.1.1	Up	DP	L2	[50]	
nve1	7500	227.1.1.1	Up	DP	L2	[75]	

VTEP-2# show running-config interface nve 1

```
interface nve1
  no shutdown
  source-interface loopback10
  member vni 5000 mcast-group 225.1.1.1
  member vni 7500 mcast-group 227.1.1.1
```

VTEP-2# show service instance vni detail

```
VSI: VSI-Ethernet7/30.1
If-index: 0x3531d001
Admin Status: Up
Oper Status: Up
Auto-configuration Mode: No
encapsulation profile vni VSI_50_TO_5000
  dot1q 50 vni 5000
Dot1q  VNI  BD
-----
50     5000  50
```

```
VSI: VSI-Ethernet7/30.2
If-index: 0x3531d002
Admin Status: Up
Oper Status: Up
Auto-configuration Mode: No
encapsulation profile vni TEST
  dot1q 100 vni 7500
Dot1q  VNI  BD
-----
100    7500  75
```

VTEP-2# show bridge-domain

```
Bridge-domain 50 (2 ports in all)
Name:: Bridge-Domain50
  Administrative State: UP                    Operational State: UP
  vni5000
  VSI-Eth7/30.1
  nve1
```

```
Bridge-domain 75 (2 ports in all)
Name:: Bridge-Domain75
  Administrative State: UP                    Operational State: UP
  vni7500
```

VSI-Eth7/30.2
nve1

VTEP-2# show mac address-table dynamic

Note: MAC table entries displayed are getting read from software.
Use the 'hardware-age' keyword to get information related to 'Age'

Legend:

* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since last seen, + - primary entry using vPC Peer-Link, E -
EVPN entry

(T) - True, (F) - False , ~~~ - use 'hardware-age' keyword to retrieve

age info

```
VLAN/BD    MAC Address      Type      age      Secure NTFY Ports/SWID.SSID.LID -----+-----
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
nve1/10.10.10.1 * 50 547f.eeec.af45 dynamic ~~~ F F VSI-Eth7/30.1 * 75 547f.eeec.af45 dynamic
~~~ F F VSI-Eth7/30.2 * 75 547f.eeec.af48 dynamic ~~~ F F nve1/10.10.10.1 VTEP-2# show ip mroute
detail IP Multicast Routing Table for VRF "default" Total number of routes: 5 Total number of
(*,G) routes: 2 Total number of (S,G) routes: 2 Total number of (*,G-prefix) routes: 1 (*,
225.1.1.1/32), uptime: 19:56:19, nve(1) ip(0) pim(0) Data Created: No VXLAN Flags VXLAN Encap
Stats: 8/748 [Packets/Bytes], 0.000 bps Incoming interface: Ethernet10/7, RPF nbr: 1.1.1.1
Outgoing interface list: (count: 1) nve1, uptime: 19:56:19, nve (10.10.10.2/32, 225.1.1.1/32),
uptime: 19:56:19, nve(0) mrib(0) pim(1) ip(0) Data Created: No Received Register stop VXLAN
Flags VXLAN Encap Stats: 9/834 [Packets/Bytes], 0.000 bps Incoming interface: loopback10, RPF
nbr: 10.10.10.2 Outgoing interface list: (count: 1) Ethernet10/7, uptime: 18:15:17, pim (*,
227.1.1.1/32), uptime: 12:57:03, nve(1) ip(0) pim(0) Data Created: No VXLAN Flags VXLAN Encap
Stats: 10/864 [Packets/Bytes], 0.000 bps Incoming interface: Ethernet10/7, RPF nbr: 1.1.1.1
Outgoing interface list: (count: 1) nve1, uptime: 12:57:03, nve (10.10.10.2/32, 227.1.1.1/32),
uptime: 12:57:03, nve(0) mrib(0) ip(0) pim(1) Data Created: No Received Register stop VXLAN
Flags VXLAN Encap Stats: 30/2648 [Packets/Bytes], 0.000 bps Incoming interface: loopback10, RPF
nbr: 10.10.10.2 Outgoing interface list: (count: 1) Ethernet10/7, uptime: 12:56:45, pim (*,
232.0.0.0/8), uptime: 18:20:36, pim(0) ip(0) Data Created: No Stats: 0/0 [Packets/Bytes], 0.000
bps Incoming interface: Null, RPF nbr: 0.0.0.0 Outgoing interface list: (count: 0) VTEP-2# show
ip arp Flags: * - Adjacencies learnt on non-active FHRP router + - Adjacencies synced via CFSOE
# - Adjacencies Throttled for Glean D - Static Adjacencies attached to down interface IP ARP
Table for context default Total number of entries: 4 Address Age MAC Address Interface
10.50.50.1 00:11:30 547f.eeec.af44 Bdi50 10.50.50.2 00:17:07 547f.eeec.af45 Bdi50
10.75.75.1 00:04:14 547f.eeec.af45 Bdi75 10.75.75.2 00:03:24 547f.eeec.af45 Bdi75 192.168.1.1
00:10:52 547f.eeec.af48 Ethernet10/7 VTEP-2# show ip route IP Route Table for VRF "default" '*'
denotes best ucast next-hop '**' denotes best mcast next-hop '[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string> 192.168.1.0/30, ubest/mbest: 1/0, attached *via
1.1.1.2, Eth10/7, [0/0], 20:30:24, direct 192.168.1.2/32, ubest/mbest: 1/0, attached *via
1.1.1.2, Eth10/7, [0/0], 20:30:24, local 10.10.10.1/32, ubest/mbest: 1/0 *via 1.1.1.1, Eth10/7,
[1/0], 20:29:48, static 10.10.10.2/32, ubest/mbest: 2/0, attached *via 10.10.10.2, Lo10, [0/0],
20:29:39, local *via 10.10.10.2, Lo10, [0/0], 20:29:39, direct 50.50.50.0/24, ubest/mbest: 1/0,
attached *via 50.50.50.51, Bdi50, [0/0], 01:22:50, direct 50.50.50.51/32, ubest/mbest: 1/0,
attached *via 50.50.50.51, Bdi50, [0/0], 01:22:50, local 75.75.75.0/24, ubest/mbest: 1/0,
attached *via 75.75.75.76, Bdi75, [0/0], 01:14:50, direct 75.75.75.76/32, ubest/mbest: 1/0,
attached *via 75.75.75.76, Bdi75, [0/0], 01:14:50, local
```

注意：所有这些输出都通过拓扑中所有主机之间的全网状流量收集。

故障排除

目前没有针对此配置的故障排除信息。