



VXLANs

This chapter describes how to identify and resolve problems that might occur when implementing Virtual Extensible Local Area Networks (VXLANs). This chapter contains the following sections:

- [Information About VXLANs, on page 1](#)
- [VXLAN Troubleshooting Commands, on page 2](#)

Information About VXLANs

Overview

A VXLAN creates LAN segments by using an overlay approach with MAC-in-UDP encapsulation and a 24-bit segment identifier in the form of a VXLAN ID. The encapsulation carries the original Layer 2 frame from the virtual machine (VM) that is encapsulated from within the Virtual Ethernet Module (VEM). Each VEM is assigned an IP address that is used as the source IP address when encapsulated MAC frames are sent over the network. You can have multiple VTEPs per VEM that are used as sources for this encapsulated traffic. The encapsulation carries the VXLAN identifier used to scope the MAC address of the payload frame. The VXLAN ID to which a VM belongs is indicated within the port profile configuration of the vNIC and is applied when the VM connects to the network. A VXLAN supports three different modes for broadcast, multicast, and MAC distribution mode transport.

For information about the VXLAN feature and how to configure it, see the [Cisco Nexus 1000V for VMware vSphere VXLAN Configuration Guide](#). For detailed information about the **show** commands mentioned in this chapter, see the [Cisco Nexus 1000V for VMware vSphere Command Reference](#).

Bridge Domains Scalability

Cisco Nexus 1000V supports a total of 4000 and 6144 bridge domains.

```
switch(config-port-prof-srv)# show resource-availability vlan
```

```
Maximum number of user VLANs supported: 4093  
Number of user VLANs created : 3968  
Total number of available user VLANs : 125  
Note: Total number of available user VLANs additionally depend on number of  
bridge-domains under usage. Please verify the usage of bridge-domains too.
```

```
VSM-DAOX(config-port-prof-srv)# show resource-availability bridge-domain  
Maximum number of bridge-domains per DVS: 6144
```

```

Number of bridge-domains currently created: 5004
Number of bridge-domains available*: 1140
* available bridge-domains do not account for created VLANs

```

VXLAN Feature Disabled

As a safety precaution, do not use the **no feature segmentation** command if there are any ports associated with a VXLAN port profile. You must remove all associations before you can disable this feature. You can use the **no feature segmentation** command to remove all the VXLAN bridge domain configurations on the Cisco Nexus 1000V.

Vempkt

Use vempkt to trace the packet path through the VEM.

- Encapsulated: Capture ingress on Seg-VEth LTL and Egress on uplink
- Decapsulated: Capture ingress on uplink and Egress on Seg-VEth LTL

VXLAN Troubleshooting Commands

VSM Show Commands

Command	Purpose
show system internal seg_bd info segment <i>segment-id</i>	Displays the ports belonging to a specific segment. See show system internal seg_bd info segment, on page 3 .
show system internal seg_bd info port vethernet	Displays the vEthernet bridge domain configuration. See show system internal seg_bd info port vethernet, on page 3 .
show system internal seg_bd info port ifindex	Displays the vEthernet bridge configuration with ifindex as an argument. See show system internal seg_bd info port ifindex, on page 3 .
show system internal seg_bd info port_count	Displays the total number of bridge domain ports. See show system internal seg_bd info port_count, on page 3 .
show system internal seg_bd info bd vxlan-home	Displays the bridge domain internal configuration. See show system internal seg_bd info bd vxlan-home, on page 3 .

Command	Purpose
<code>show system internal seg_bd info port</code>	Displays the VXLAN vEthernet information. See show system internal seg_bd info port, on page 3 .

show system internal seg_bd info segment

```
switch(config)# show system internal seg_bd info segment 10000
Bridge-domain: A
Port Count: 11
Veth1
Veth2
Veth3
```

show system internal seg_bd info port vethernet

```
show system internal seg_bd info port vethernet 1
Bridge-domain: A
segment_id = 10000
Group IP: 225.1.1.1
```

show system internal seg_bd info port ifindex

```
switch(config)# show system internal seg_bd info port ifindex 0x1c000050
Bridge-domain: A
segment_id = 10000
Group IP: 225.1.1.1
```

show system internal seg_bd info port_count

```
switch(config)# show system internal seg_bd info port_count
Number of ports: 11
```

show system internal seg_bd info bd vxlan-home

```
switch(config)# show system internal seg_bd info bd vxlan-home

Bridge-domain vxlan-home (2 ports in all)
Segment ID: 5555 (Manual/Active)
Group IP: 235.5.5.5
State: UP Mac learning: Enabled
is_bd_created: Yes
current state: SEG_BD_FSM_ST_READY
pending_delete: 0
port_count: 2
action: 4
hwbd: 28
pa_count: 0
Veth2, Veth5
switch(config)#
```

show system internal seg_bd info port

```
switch# show system internal seg_bd info port
if_index = <0x1c000010>
Bridge-domain vxlan-pepsi
rid = 216172786878513168
swbd = 4098
```

```

if_index = <0x1c000040>
Bridge-domain vxlan-pepsi
rid = 216172786878513216
swbd = 4098

```

```
switch#
```

BGP Show Commands

The following table describes the BGP show commands. For detailed information about these commands, see the [Cisco Nexus 1000V Command Reference](#).

For information about how to configure BGP and peer templates, see the [Cisco Nexus 1000V for VMware vSphere VXLAN Configuration Guide](#).

Command	Purpose
show bgp session	Displays the BGP sessions.
show bgp l2vpn evpn	Displays the VTEPs that are learned through the BGP.
show bgp l2vpn evpn rd	Displays the detailed output for a specific segment ID or RD.
show bgp convergence	Displays the BGP convergence time.
show bgp l2vpn evpn evi all VTEP	Displays the VTEP list for a specific VXLAN segment ID or all segments.
show bridge-domain VTEPs	Displays the bridge domain-to-VTEP mappings that are maintained by the VSM and are pushed to all VEMs.
show bgp l2vpn evpn summary	Displays the BGP evpn summary.
show bgp l2vpn evpn neighbors	Displays the detailed state for a neighbor.
show bgp internal evi	Displays the detailed state for a VXLAN segment.
show bgp event-history msgs	Displays various message logs of BGP.
show bgp event-history events	Displays event logs.

VEM Show Commands

Command	Purpose
vemcmd show port segments	Displays VXLAN vEthernet programming.
vemcmd show vxlan interfaces	Displays the VXLAN encapsulated interfaces.
vemcmd show port vlans	Checks the port programming and CBL state for the bridge domain.

Command	Purpose
vemcmd show bd	Displays the bridge domain segment ID, group, or list of ports.
vemcmd show bd bd-name <i>bd-name-string</i>	Displays one segment bridge domain.
vemcmd show l2 all	Displays the remote IP being learned.
vemcmd show l2 bd-name <i>bd-name-string</i>	Displays the layer 2 table for one segment bridge domain.
vemcmd show arp all	Displays the IP-MAC mapping for the outer encapsulated header.

VXLAN Gateway Commands



Note Starting with Release 5.2(1)SV3(1.15), Cisco Nexus 1000V does not support the VXLAN Gateway feature.

Command	Purpose
show module vem	Displays VXLAN Gateway information that is attached to the VSM.
attach vem	Displays VXLAN Gateway information that is not attached to the VSM.
vemcmd show vxlan-gw-mappings	Displays VXLAN Gateway mappings. See vemcmd show vxlan-gw-mappings, on page 6 .
vemcmd show vxlan-stats	Displays VXLAN Gateway statistics. See vemcmd show vxlan-stats, on page 6 .
vemlog show all	Displays the VXLAN Gateway packet path.
show bridge-domain	Displays the bridge-domain configuration on the VSM.
show bridge-domain VTEPs	Displays the bridge-domain VTEPs on the VSM.
show bridge-domain mapping	Displays the VLAN-VXLAN mappings programmed on the VSM.
show module vteps	Displays the interfaces on the VSM.
show bridge-domain vteps	Displays the bridge domain-to-VTEP mappings that are maintained by the VSM and are pushed to all VEMs
show int switchport	Displays the port configuration on the VSM.
show bridge-domain segment-cisco VTEPs	Displays the VTEP distribution on the VSM.

Command	Purpose
show bridge-domain mac	Displays VXLAN mac-distribution.
show platform fwm	Displays the VTEPs information.

vemcmd show vxlan-gw-mappings

```
VXGW-switch(vem-attach)# vemcmd show vxlan-gw-mappings
VLAN Segment NumProbes State
-----
1821 9001 3 Active
1822 9002 3 Active
Linux(debug)#
Linux(debug)#
Linux(debug)# vemcmd show vxlan
LTL VSM Port IP Seconds since Last Vem Port
Netmask IGMP Query Received
Gateway
(* = IGMP Join Interface/Designated VTEP)
-----
20 Veth7 17.17.19.111 33 vxlannic0 *
255.255.255.0
17.17.19.1
```

vemcmd show vxlan-stats

```
switch(vem-attach)# vemcmd show vxlan-stats
LTL Ucast Mcast/Repl Ucast Mcast Total
Encaps Encaps Decaps Decaps Drops
17 8717 173 8334 0 242
switch(vem-attach)#
switch(vem-attach)# vemcmd show vxlan-stats ltl 17
VXLAN Port Stats for LTL 17
Unicast Encapsulations: 8756
Multicast Encapsulations/HeadEnd Replications: 173
Unicast Decapsulations: 8372
Multicast Decapsulations: 0
IP Pre-fragmentations: 0
TSO Processed Packets: 0
ICMP Pkt Too Big msgs from upstream: 0
ICMP Pkt Too Big msgs sent to VM: 0
Packets generated by Head End Replication: 172
```

VEM Packet Path Debugging Commands

Use the following commands to debug VXLAN traffic from a VM on VEM1 to a VM on VEM2.

VEM	Command	Purpose
VEM 1	vempkt capture ingress ltl vxlan_veth	Verifies that packets are coming into the switch from the segment vEthernet.
VEM1	vemlog debug sflisp all vemlog debug sfvnsegment all	Verifies VXLAN encapsulation.

VEM	Command	Purpose
VEM1	vemcmd show l2 bd-name segbdname	Verifies that the remote IP address is learned. If the remote IP is not learned, packets are sent multicast encapsulated.
VEM1	vemcmd show vxlan-encap ltl vempkt capture egress ltl uplink	Verifies that encapsulated packets go out on an uplink. Use the vemcmd show vxlan-encap ltl command to find out which uplink is being used.
VEM1	vemcmd show vxlan-stats all vemcmd show vxlan-stats ltl veth/vxlanVTEP	Displays statistics that can be used to find information about any failures.
VEM2	vempkt capture ingress ltl uplink	Verifies encapsulated packets are arriving on the uplink.
VEM2	vemlog debug sflisp all vemlog debug sfvsegment all	Verifies VXLAN decapsulation.
VEM2	vempkt capture egress ltl vxlan_veth	Verifies that the decapsulated packets go out on a VXLAN vEthernet interface.
VEM2	vemcmd show vxlan-stats all vemcmd show vxlan-stats ltl veth/vxlanVTEP	Displays statistics that can be used to find information about any failures.

VEM Multicast Debugging Commands

Commands	Purpose
vemcmd show igmp vxlan_transport_vlan detail	Verifies the IGMP state on the VEM. Note This command does not show any output for the segment multicast groups. To save multicast table space, segment groups are not tracked by IGMP snooping on the VEM.
vemcmd show vxlan interfaces	Verifies that the IGMP queries are being received.
vempkt capture ingress ltl first_vxlan_VTEP_ltl	Verifies that the VMware stack is sending joins.
vempkt capture egress ltl uplink_ltl	Verifies that the joins are being sent out to the upstream switch.

VXLAN Data Path Debugging

vemlog Debugging Commands

Command	Purpose
<code>vemlog debug sfbid all</code>	Debugs the bridge domain setup or configuration.
<code>vemlog debug sfporttable all</code>	Debugs the port configuration, CBL, vEthernet LTL pinning.
<code>vemlog debug sfvnsegment all</code>	Debugs the encapsulated/decapsulated setup.
<code>vemlog debug sflisp all</code>	Debugs for actual packet editing, VXLAN interface handling, and multicast handling.
<code>echo "debug dpa_allplatform all" > /tmp/dpafifo</code>	Debugs the multicast joins or leaves on the DPA socket.
<code>echo "debug sfl2agent all" > /tmp/dpafifo</code>	Debugs the bridge domain configuration.
<code>echo "debug sfportagent all" > /tmp/dpafifo</code>	Debugs the port configuration.
<code>echo "debug sfportl2lisp_cache all" > /tmp/dpafifo</code>	Debugs the hitless reconnect (HR) for capability l2-lisp.
<code>echo "debug sfpixmagent all" > /tmp/dpafifo</code>	Debugs CBL programming.
<code>echo "debug sfvxlanagent all" > /tmp/dpafifo</code>	Debugs a VXLAN agent that interacts with the VSM.

VEM Statistics Commands

Command	Purpose
<code>vemcmd show vxlan-stats</code>	Displays a summary of per-port statistics.
<code>vemcmd show vxlan-stats ltl <i>vxlan_VTEP_ltl</i></code>	Displays detailed per-port statistics for VXLAN VTEP.
<code>vemcmd show vxlan-stats ltl <i>vxlan_veth_ltl</i></code>	Displays detailed per-port statistics for the vEthernet interface in a VXLAN.
<code>vemcmd show vxlan-stats ltl <i>vxlan_VTEP_ltl</i> bd-all</code>	Displays detailed per-port-per-bridge domain statistics for a VXLAN VTEP for all bridge domains.
<code>vemcmd show vxlan-stats ltl <i>vxlan_VTEP_ltl</i> bd-name <i>bd-name</i></code>	Displays detailed per-port-per-bridge domain statistics for a VXLAN VTEP for the specified bridge domain.
<code>vemcmd show vxlan-encap ltl <i>vxlan_veth_ltl</i></code>	Displays which VXLAN VTEP is used for encapsulation and subsequent pinning to the uplink port channel for static MAC addresses learned on port.

Command	Purpose
vemcmd show vxlan-encap mac <i>vxlan_vm_mac</i>	Displays which VXLAN VTEP is used for encapsulation and subsequent pinning to the uplink port channel.

