



## Layer 2 Switching

---

This chapter describes how to identify and resolve problems that relate to Layer 2 switching and includes the following sections:

- [Information About Layer 2 Ethernet Switching, page 10-1](#)
- [Port Model, page 10-1](#)
- [Layer 2 Switching Problems, page 10-4](#)
- [Layer 2 Switching Troubleshooting Commands, page 10-6](#)

### Information About Layer 2 Ethernet Switching

The Cisco Nexus1000VE is a distributed Layer 2 virtual switch that extends across many virtualized hosts.

It consists of two components:

- The Virtual Supervisor Module (VSM), which is also known as the control plane (CP). The VSM acts as the supervisor and contains the Cisco CLI, configuration, and high-level features.
- The Virtual Services Engine (VSE), which is also known as the data plane (DP). The VSE acts as a line card and runs as a VM in each virtualized server to handle packet forwarding and other localized functions.

### Port Model


This section includes the following topics:

- [Viewing Ports from the VSE, page 10-2](#)
- [Viewing Ports from the VSM, page 10-3](#)

## Viewing Ports from the VSE

The Cisco Nexus1000VE differentiates between virtual and physical ports on each of the VSEs. Figure 10-1 shows how ports on the Cisco Nexus1000VE switch are bound to physical and virtual VMware ports within a VSE.

Figure 10-1 VSE View of Ports



On the virtual side of the switch, three layers of ports are mapped together:

- **Virtual NICs**—Three types of Virtual NICs are in VMware. The virtual NIC (vnic) is part of the VM and represents the physical port of the host that is plugged into the switch. The virtual kernel NIC (VTEP) is used by the hypervisor for management, VMotion, iSCSI, network file system (NFS), and other network access needed by the kernel. This interface carries the IP address of the hypervisor itself and is also bound to a virtual Ethernet port. The vswif (not shown) appears only in CoS-based systems and is used as the VMware management port. Each type maps to a virtual Ethernet port within the Cisco Nexus1000VE.
- **Virtual Ethernet Ports (VEth)**—A vEth port is a port on the Cisco Nexus 1000V. The Cisco Nexus 1000V has a flat space of vEth ports 0..N. The virtual cable plugs into these vEth ports that are moved to the host running the VM.

Virtual Ethernet ports are assigned to port groups.

- **Local Virtual Ethernet Ports (lveth)**—Each host has a number of local vEth ports. These ports are dynamically selected for vEth ports that are needed on the host.

These local ports do not move and are addressable by the module/port number method.


On the physical side of the switch, from bottom to top, is the following:

- Each physical NIC in VMware is represented by an interface called a vmnic. The vmnic number is allocated during VMware installation, or when a new physical NIC is installed, and remains the same for the life of the host.

## Viewing Ports from the VSM

Figure 10-2 shows the VSM view ports.

**Figure 10-2** VSM View of Ports



## Port Types

The following types of ports are available:

- vEths can be associated with any one of the following:
  - VNICs of a Virtual Machine on the ESX host.
  - VTEPs of the ESX Host
  - VSWIFs of an ESX COS Host.
- Eths (physical Ethernet interfaces)—Correspond to the outside-trunk interface of the VSEs.

For more information about Layer 2 switching, see the *Cisco Nexus 1000VE Layer 2 Switching Configuration Guide*.

## Layer 2 Switching Problems

This section describes how to troubleshoot Layer 2 problems and lists troubleshooting commands. This section includes the following topics:

- [Verifying a Connection Between VSE Ports, page 10-4](#)
- [Verifying a Connection Between VSEs, page 10-4](#)
- [Isolating Traffic Interruptions, page 10-5](#)

### Verifying a Connection Between VSE Ports

You can verify a connection between two vEth ports on a VSE.

- Step 1** View the state of the VLANs associated with the port. If the VLAN associated with a port is not active, the port may be down. In this case, you must create the VLAN and activate it.

```
switch# show vlan vlan-id
```

- Step 2** View the state of the ports on the VSM.

```
switch# show interface brief
```

- Step 3** Display the ports that are present on the VSE, their local interface indices, VLAN, type (physical or virtual), port mode and port name.

```
switch# module vse module-number execute vemcmd show port
```

The key things to look for in the output are as follows:

- State of the port.
- CBL.
- Mode.
- Attached device name.
- The LTL of the port that you are trying to troubleshoot. It will help you to identify the interface quickly in other VSE commands where the interface name is not displayed.
- Make sure that the state of the port is up. If not, verify the configuration of the port on the VSM.

- Step 4** View the VLANs and port lists on a particular VSE.

```
switch# module vse module-number execute vemcmd show bd
```

If you are trying to verify that a port belongs to a particular VLAN, make sure that you see the port name or LTL in the port list of that VLAN.

### Verifying a Connection Between VSEs

You can verify a connection between vEth ports on two separate VSEs.

- Step 1** Log in to the upstream switch and make sure that the port is configured to allow the VLAN that you are looking for.

```
switch# show running-config interface gigabitEthernet 1/38
Building configuration...
```

```

Current configuration : 161 bytes
!
interface GigabitEthernet1/38
  description Srvr-100:vmnic1
  switchport
  switchport trunk allowed vlan 1,60-69,231-233
  switchport mode trunk
end

```

As this output shows, VLANs 1,60-69, 231-233 are allowed on the port. If a particular VLAN is not in the allowed VLAN list, make sure to add it to the allowed VLAN list of the port.

## Isolating Traffic Interruptions

You can isolate the cause for no traffic passing across VMs on different VSEs.

**Step 1** Inside the VM, verify that the Ethernet interface is up.

**ifconfig -a**

If not, delete that NIC from the VM, and add another NIC.

**Step 2** Using any sniffer tool, verify that ARP requests and responses are received on the VM interface.

**Step 3** On the upstream switch, look for the association between the IP and MAC address:

**debug arp**

**show arp**

**Example:**

```

switch# debug arp
ARP packet debugging is on
11w4d: RARP: Rcvd RARP req for 0050.56b7.3031
11w4d: RARP: Rcvd RARP req for 0050.56b7.3031
11w4d: RARP: Rcvd RARP req for 0050.56b7.4d35
11w4d: RARP: Rcvd RARP req for 0050.56b7.52f4
11w4d: IP ARP: rcvd req src 10.78.1.123 0050.564f.3586, dst 10.78.1.24 Vlan3002
11w4d: RARP: Rcvd RARP req for 0050.56b7.3031
switch#

```

**Example:**

```

switch# show arp

```

Protocol	Address	Age (min)	Hardware Addr	Type	Interface
Internet	10.78.1.72	-	001a.6464.2008	ARPA	
Internet	7.114.1.100	-	0011.bcac.6c00	ARPA	Vlan140
Internet	41.0.0.1	-	0011.bcac.6c00	ARPA	Vlan410
Internet	7.61.5.1	-	0011.bcac.6c00	ARPA	Vlan1161
Internet	10.78.1.5	-	0011.bcac.6c00	ARPA	Vlan3002
Internet	7.70.1.1	-	0011.bcac.6c00	ARPA	Vlan700
Internet	7.70.3.1	-	0011.bcac.6c00	ARPA	Vlan703
Internet	7.70.4.1	-	0011.bcac.6c00	ARPA	Vlan704
Internet	10.78.1.1	0	0011.bc7c.9c0a	ARPA	Vlan3002
Internet	10.78.1.15	0	0050.56b7.52f4	ARPA	Vlan3002
Internet	10.78.1.123	0	0050.564f.3586	ARPA	Vlan3002

**Step 4** You have completed this procedure.

## Layer 2 Switching Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to the Layer 2 MAC address configuration.

Command	Purpose
<b>show mac address-table</b>	Displays the MAC address table to verify all MAC addresses on all VSEs controlled by the VSM. See <a href="#">Example 10-1 on page 10-7</a> .
<b>show mac address-table module</b> <i>module-number</i>	Displays all the MAC addresses on the specified VSE.
<b>show mac address-table static</b> <i>HHHH.WWWW.HHHH</i>	Displays the MAC address table static entries. See <a href="#">Example 10-2 on page 10-8</a> .
<b>show mac address-table address</b> <i>HHHH.WWWW.HHHH</i>	Displays the interface on which the MAC address specified is learned or configured. <ul style="list-style-type: none"> <li>For dynamic MAC addresses, if the same MAC address appears on multiple interfaces, each of them is displayed separately.</li> <li>For static MAC addresses, if the same MAC address appears on multiple interfaces, only the entry on the configured interface is displayed.</li> </ul>
<b>show mac address-table static   inc veth</b>	Displays the static MAC address of vEthernet interfaces in case a VSE physical port learns a dynamic MAC address and the packet source is in another VSE on the same VSM. See <a href="#">Example 10-3 on page 10-8</a> .
<b>show running-config vlan</b> <i>vlan-id</i>	Displays VLAN information in the running configuration.
<b>show vlan</b> [ <b>all-ports</b>   <b>brief</b>   <b>id</b> <i>vlan-id</i>   <b>name</b> <i>name</i>   <b>dot1q tag native</b> ]	Displays VLAN information as specified. See <a href="#">Example 10-4 on page 10-8</a> .
<b>show vlan summary</b>	Displays a summary of VLAN information.
<b>show interface brief</b>	Displays a table of interface states. See <a href="#">Example 10-5 on page 10-9</a> .
<b>module vse</b> <i>module-number</i> <b>execute vemcmd</b> <b>show port</b>	On the VSE, displays the port state on a particular VSE.  This command can only be used from the VSE. See <a href="#">Example 10-6 on page 10-9</a> .
<b>module vse</b> <i>module-number</i> <b>execute vemcmd</b> <b>show bd</b>	For the specified VSE, displays its VLANs and their port lists. See <a href="#">Example 10-7 on page 10-10</a> .

Command	Purpose
<b>module vse <i>module-number</i> execute vemcmd show trunk</b>	For the specified VSE, displays the VLAN state on a trunk port. <ul style="list-style-type: none"> <li>If a VLAN is forwarding (active) on a port, its CBL state should be 1.</li> <li>If a VLAN is blocked, its CBL state is 0.</li> </ul> See <a href="#">Example 10-8 on page 10-11</a> .
<b>module vse <i>module-number</i> execute vemcmd show l2 <i>vlan-id</i></b>	For the specified VSE, displays the VLAN forwarding table for a specified VLAN. See <a href="#">Example 10-9 on page 10-11</a> .
<b>show interface <i>interface_id</i> mac</b>	Displays the MAC addresses and the burn-in MAC address for an interface.

### Example 10-1 show mac address-table Command



**Note** The Cisco Nexus 1000V MAC address table does not display multicast MAC addresses.



**Tip** The “Module” indicates the VSE on which this MAC address is seen.

The “N1KV Internal Port” refers to an internal port created on the VSE. This port is used for control and management of the VSE and is not used for forwarding packets.

```
switch# show mac address-table
VLAN      MAC Address           Type    Age    Port                               Mod
-----+-----+-----+-----+-----+-----
1         0002.3d23.7802       static  0      N1KV Internal Port               3
1         0002.3d33.7802       static  0      N1KV Internal Port               3
1         0002.3d43.7802       static  0      N1KV Internal Port               3
1         0002.3d63.7802       static  0      N1KV Internal Port               3
1         0002.3d83.7802       static  0      N1KV Internal Port               3
222      0050.56b8.7584       static  0      Veth2                             3
222      d48c.b5bc.fe01       dynamic 0      Eth3/1                             3
223      0050.56b8.0375       static  0      Veth1                             3
3968     0002.3d83.7802       static  0      N1KV Internal Port               3
3970     0002.3d83.7802       static  0      N1KV Internal Port               3
3971     0002.3d83.7802       static  0      N1KV Internal Port               3
3972     0002.3d83.7802       static  0      N1KV Internal Port               3
1         0002.3d23.7803       static  0      N1KV Internal Port               4
1         0002.3d33.7803       static  0      N1KV Internal Port               4
1         0002.3d43.7803       static  0      N1KV Internal Port               4
1         0002.3d63.7803       static  0      N1KV Internal Port               4
1         0002.3d83.7803       static  0      N1KV Internal Port               4
222      0050.56b8.8ce8       static  0      Veth3                             4
223      0050.56b8.99b6       static  0      Veth4                             4
3968     0002.3d83.7803       static  0      N1KV Internal Port               4
3970     0002.3d83.7803       static  0      N1KV Internal Port               4
3971     0002.3d83.7803       static  0      N1KV Internal Port               4
3972     0002.3d83.7803       static  0      N1KV Internal Port               4
Total MAC Addresses: 23
```

**Example 10-2 show mac address-table address Command**

**Tip** This command shows all interfaces on which a MAC is learned dynamically. In this example, the same MAC appears on Eth3/1 and Eth4/1.

```
switch# show mac address-table address 0050.568d.5a3f
VLAN      MAC Address      Type    Age    Port      Module
-----+-----+-----+-----+-----+-----
342      0050.568d.5a3f  dynamic 0      Eth3/3    3
342      0050.568d.5a3f  dynamic 0      Eth4/3    4
Total MAC Addresses: 1
switch#
```

**Example 10-3 show mac address-table static | inc veth Command**

```
switch# show mac address-table static | inc veth
460      0050.5678.ed16  static 0      Veth2     3
460      0050.567b.1864  static 0      Veth1     4
switch#
```

**Example 10-4 show vlan Command**

**Tip** This command shows the state of each VLAN created on the VSM.

```
switch# show vlan
VLAN Name                Status  Ports
-----+-----+-----+-----
1    default                active  Eth3/1, Eth4/1
110  VLAN0110               active
111  VLAN0111               active
112  VLAN0112               active
113  VLAN0113               active
114  VLAN0114               active
115  VLAN0115               active
116  VLAN0116               active
117  VLAN0117               active
118  VLAN0118               active
119  VLAN0119               active
800  VLAN0800               active
801  VLAN0801               active
802  VLAN0802               active
803  VLAN0803               active
804  VLAN0804               active
805  VLAN0805               active
806  VLAN0806               active
807  VLAN0807               active
808  VLAN0808               active
809  VLAN0809               active
810  VLAN0810               active
811  VLAN0811               active
812  VLAN0812               active
813  VLAN0813               active
814  VLAN0814               active
815  VLAN0815               active
816  VLAN0816               active
```



```

817 VLAN0817          active
818 VLAN0818          active
819 VLAN0819          active
820 VLAN0820          active
VLAN Name            Status    Ports
-----
-----

Remote SPAN VLANs
-----

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

```

**Example 10-5 show interface brief Command**

```

switch# show interface brief
-----
Port      VRF      Status IP Address          Speed  MTU
-----
mgmt0    --      up    172.23.232.163      1000  1500
-----

Ethernet  VLAN    Type Mode    Status Reason          Speed  Port
Interface                                Speed Reason          Speed  Ch #
-----
Eth3/1    1       eth trunk up    none          10G
Eth4/1    1       eth trunk up    none          10G
-----

Vethernet VLAN/   Type Mode    Status Reason          MTU  Module
          Segment
-----
Veth1     223    virt access up    none          1500 3
Veth2     222    virt access up    none          1500 3
Veth3     222    virt access up    none          1500 4
Veth4     223    virt access up    none          1500 4
-----

Port      VRF      Status IP Address          Speed  MTU
-----
control0  --      up    --                  1000  1500
-----

NOTE : * Denotes ports on modules which are currently offline on VSM

```

**Example 10-6 module vse module-number execute vemcmd show port Command**

**Tip** Look for the state of the port.

```

siwtch# module vse 3 execute vemcm show port
LTL  VSM Port  Admin Link  State  PC-LTL  SGID          Vem Port  Type  ORG
svcpath Owner
  21   Eth3/1   UP   UP   F/B*    0           eth1      0
0 dpdk-outside
  53   Veth2    UP   UP   FWD     0           test-vm1.eth1  0
0 test-vm1
  54   Veth1    UP   UP   FWD     0           test-vm2.eth1  0
0 test-vm2

```

\* F/B: Port is BLOCKED on some of the vlans.  
 One or more vlans are either not created or  
 not in the list of allowed vlans for this port.  
 Please run "vemcmd show port vlans" to see the details.

### Example 10-7 module vse module-number execute vemcmd show bd Command



**Tip** If a port belongs to a particular VLAN, the port name or LTL should be in the port list for the VLAN.

```
switch# module vse 3 execute vemcmd show bd
BD 1, vdc 1, vlan 1, swbd 1, table-id 0, 1 ports, ""
Forward type: L2
Portlist:
    12 _l24

BD 2, vdc 1, vlan 3972, swbd 3972, table-id 0, 0 ports, ""
Forward type: L2
Portlist:
BD 3, vdc 1, vlan 3970, swbd 3970, table-id 0, 0 ports, ""
Forward type: L2
Portlist:
BD 4, vdc 1, vlan 3968, swbd 3968, table-id 0, 1 ports, ""
Forward type: L2
Portlist:
    11 _l23

BD 5, vdc 1, vlan 3971, swbd 3971, table-id 0, 1 ports, ""
Forward type: L2
Portlist:
    15 _l27

BD 6, vdc 1, vlan 222, swbd 222, table-id 0, 2 ports, ""
Forward type: L2
Portlist:
    21 eth1
    53 test-vm1.eth1

BD 7, vdc 1, vlan 220, swbd 220, table-id 0, 1 ports, ""
Forward type: L2
Portlist:
    21 eth1

BD 8, vdc 1, vlan 221, swbd 221, table-id 0, 1 ports, ""
Forward type: L2
Portlist:
    21 eth1

BD 9, vdc 1, vlan 223, swbd 223, table-id 0, 2 ports, ""
Forward type: L2
Portlist:
    21 eth1
    54 test-vm2.eth1
```

**Example 10-8** *module vse module-number execute vemcmd show trunk Command*

**Tip** If a VLAN is active on a port, its CBL state should be 1.  
If a VLAN is blocked, its CBL state is 0.

```
switch# module vse 3 execute vemcmd show trunk
Trunk port 6 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(3972) cbl 1, vlan(3970) cbl 1, vlan(3968) cbl 1, vlan(3971) cbl 1,
vlan(222) cbl 1, vlan(220) cbl 1, vlan(221) cbl 1, vlan(223) cbl 1, vlan(224) cbl 1,
vlan(225) cbl 1, vlan(226) cbl 1, vlan(227) cbl 1, vlan(228) cbl 1, vlan(229) cbl 1,
Trunk port 16 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(3972) cbl 1, vlan(3970) cbl 1, vlan(3968) cbl 1, vlan(3971) cbl 1,
vlan(222) cbl 1, vlan(220) cbl 1, vlan(221) cbl 1, vlan(223) cbl 1, vlan(224) cbl 1,
vlan(225) cbl 1, vlan(226) cbl 1, vlan(227) cbl 1, vlan(228) cbl 1, vlan(229) cbl 1,
Trunk port 21 native_vlan 1 CBL 0
vlan(222) cbl 1, vlan(220) cbl 1, vlan(221) cbl 1, vlan(223) cbl 1, vlan(224) cbl 1,
vlan(225) cbl 1, vlan(226) cbl 1, vlan(227) cbl 1, vlan(228) cbl 1, vlan(229) cbl 1,
switch#
switch# module vse 3 execute vemcmd show l2
switch# module vse 3 execute vemcmd show l2 222
Bridge domain 6 brtmax 4096, brtcnt 2, timeout 300
VLAN 222, swbd 222, ""
Flags: P - PVLAN S - Secure D - Drop R - Router-mac
      Type      MAC Address  LTL  timeout  Flags  PVLAN
Dynamic  d4:8c:b5:bc:fe:01  21    1        1
Static   00:50:56:b8:75:84   53    0        0
```

**Example 10-9** *module vse module-number execute vemcmd show l2 Command*

```
~ # module vse 5 execute vemcmd show l2
Bridge domain 115 brtmax 1024, brtcnt 2, timeout 300
Dynamic MAC 00:50:56:bb:49:d9 LTL 16 timeout 0
Dynamic MAC 00:02:3d:42:e3:03 LTL 10 timeout 0
```

## Limitations and Restrictions

A syslog is generated if one of the following configurations exists when you try to disable automatic static MAC learning for MS-NLB because they do not support this feature:

- PVLAN port
- Ports configured with unknown unicast flood blocking (UUFB)

## Disabling Automatic Static MAC Learning on a vEthernet Interface

You must disable automatic static MAC learning before you can successfully configure NLB on a vEthernet (vEth) interface.

In interface configuration mode enter the following commands:

```
switch(config)# int veth 1
switch(config-if)# no mac auto-static-learn
```

In port profile configuration mode enter the following commands:

```
switch(config)# port-profile type vethernet ms-nlb
```

```
switch(config-port-prof)# no mac auto-static-learn
```

## Checking Status on a VSM

If the NLB unicast mode configuration does not function, check the status of the Virtual Supervisor Module (VSM).

Confirm that the **no mac auto-static-learn** command is listed in the vEth and/or port profile configurations.

**Step 1** In interface configuration mode, generate the VSM status.

```
switch(config-if)# show running-config int veth1
interface Vethernet1
  inherit port-profile vm59
  description Fedora117, Network Adapter 2
  no mac auto-static-learn
  vmware dvport 32 dvswitch uuid "ea 5c 3b 50 cd 00 9f 55-41 a3 2d 61 84 9e 0e c4"
```

**Step 2** In port profile configuration mode, generate the VSM status.

```
switch(config-if)# show running-config port-profile ms-nlb
port-profile type vethernet ms-nlb
  vmware port-group
  switchport mode access
  switchport access vlan 59
  no mac auto-static-learn
  no shutdown
  state enabled
```

## Checking the Status on a VSE

If the NLB unicast mode configuration does not function, check the status of the Virtual Ethernet Module (VSE). Check the following:

- Confirm that the MS-NLB vEths are disabled.
- Confirm that the MS-NLB shared-MAC (starting with 02:BF) is not listed in the Layer 2 (L2) MAC table.

**Step 1** Generate the VSE status.

```
~ # vemcmd show port auto-smac-learning
LTL   VSM Port  Auto Static MAC Learning
 49   Veth4    DISABLED
 50   Veth5    DISABLED
 51   Veth6    DISABLED
```

**Step 2** Generate the Layer 2 MAC address table for VLAN 59.

```
~ # vemcmd show 12 59
Bridge domain 15 brtmax 4096, brtcnt 6, timeout 300
VLAN 59, swbd 59, ""
Flags: P - PVLAN S - Secure D - Drop
      Type          MAC Address  LTL  timeout  Flags  PVLAN
Dynamic 00:15:5d:b4:d7:02 305    4
Dynamic 00:15:5d:b4:d7:04 305    25
Dynamic 00:50:56:b3:00:96  51     4
```

```
Dynamic 00:50:56:b3:00:94 305 5
Dynamic 00:0b:45:b6:e4:00 305 5
Dynamic 00:00:5e:00:01:0a 51 0
```

