



Layer 2 Switching

This chapter describes how to identify and resolve problems that relate to Layer 2 switching.

Information About Layer 2 Ethernet Switching

The Cisco Nexus 1000V provides a distributed, Layer 2 virtual switch that extends across many virtualized hosts.

It consists of two components:

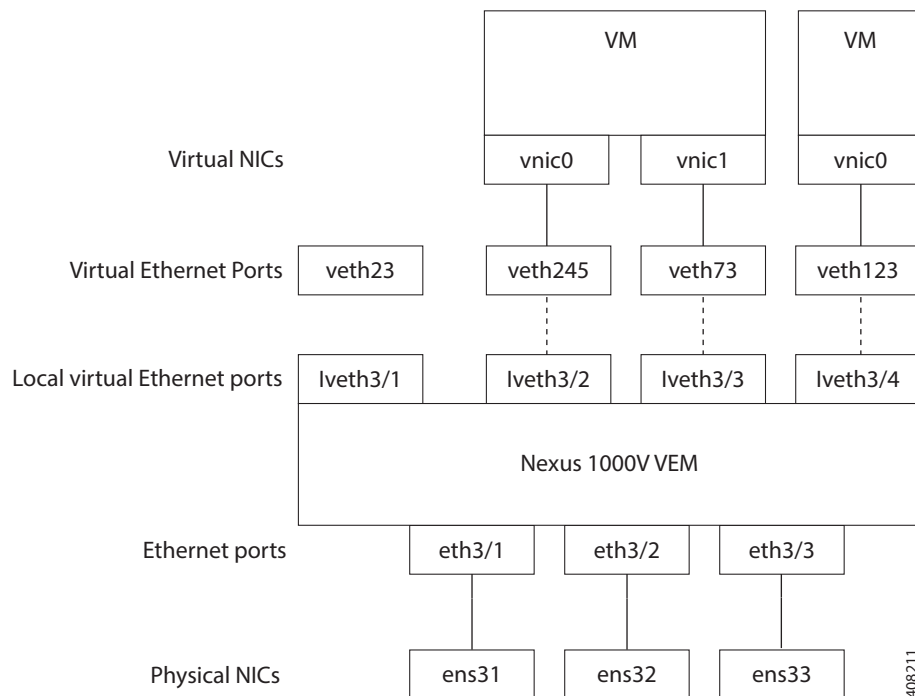
- Virtual Supervisor Module (VSM), which is also known as the control plane (CP), acts as the supervisor and contains the Cisco CLI, configuration, and high-level features.
- Virtual Ethernet Module (VEM), which is also known as the data plane (DP), acts as a line card and runs in each virtualized server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

Viewing Ports from the VEM

The Cisco Nexus 1000V differentiates between virtual and physical ports on each of the VEMs.

Figure 8-1 shows how ports on the Cisco Nexus 1000V switch are bound to physical and virtual ports within a VEM.

Figure 8-1 VEM View of Ports



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

- **Virtual NICs**—There are two types of Virtual NICs. The virtual NIC (vnic) is part of the VM and represents the physical port of the host that is plugged into the switch. Internal NICs are used by the hypervisor for internal purposes. Each type maps to a vEth port within the Cisco Nexus 1000V.
- **Virtual Ethernet Ports (vEth)**—A vEth port is a port on the Cisco Nexus 1000V distributed virtual switch. 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 that is running the VM.

vEth 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 you can address them by the module-port number method.


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

Each uplink port on the host represents a physical interface. Each physical port that is added to the Cisco Nexus 1000V switch appears as a physical Ethernet port, just as it would on a hardware-based switch.

Viewing Ports from the VSM

Figure 8-2 shows the VSM view of the ports.

Figure 8-2 VSM View of Ports



277344

Port Types

The following types of ports are available:

- vEths (virtual Ethernet interfaces) can be associated with any one of the following:
 - vNICs of a VM on the hypervisor.
 - Internal NICs on the hypervisor.
- eths (physical Ethernet interfaces)—Correspond to the physical NICs on the hypervisor.
- Po (port channel interfaces)—The physical NICs of a hypervisor can be bundled into a logical interface. This logical bundle is referred to as a port channel interface.

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

Problems with Layer 2 Switching

This section describes how to troubleshoot Layer 2 problems and lists troubleshooting commands.

Verifying a Connection Between VEM Ports

-
- Step 1** View the state of the VLANs associated with the port by entering the **show vlan** command on the VSM. If the VLAN associated with a port is not active, the port might be down. In this case, you must create the VLAN and activate it.
- Step 2** To see the state of the port on the VSM, enter the **show interface brief** command.
- Step 3** Display the ports that are present on the VEM, their local interface indices, VLAN, type (physical or virtual), CBL state, port mode, and port name by entering the **module vem module-number execute vemcmd show port** command.

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 identify the interface quickly in other VEM 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 their port lists on a particular VEM by entering the **module vem module-number execute vemcmd show bd** command.

```
switch# module vem 5 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 VEMs

-
- Step 1** Check if the VLAN associated with the port is created on the VSM by entering the **show vlan** command.
- Step 2** Check if the ports are up in the VSM by entering the **show interface brief** command.
- Step 3** Check if the CBL state of the two ports is set to the value of 1 for forwarding (active) by entering the **module vem 3 execute vemcmd show port** command on the VEM.
- Step 4** Check if the two vEth ports are listed in the flood list of the VLAN to which they are trying to communicate by entering the **module vem 3 execute vemcmd show bd** command on the VEM.
- Step 5** Verify that the uplink switch to which the VEMs are connected is carrying the VLAN to which the ports belong.
- Step 6** Find the port on the upstream switch to which the physical NIC (that is supposed to be carrying the VLAN) on the VEM is connected to.

```
switch# show cdp neighbors
```

```
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute
```

Device ID	Local Intrfce	Hldtme	Capability	Platform	Port ID
swordfish-6k-2	Eth5/2	168	R S I	WS-C6506-E	Gig1/38

The PNIC (Eth 5/2) is connected to swordfish-6k-2 on port Gig1/38.

- Step 7** Log in to the upstream switch and make sure 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 to 69 and 231 to 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

- Step 1** In the output of the **show port-profile name** command, verify the following information:

- The control and packet VLANs that you configured are present (in the example, these VLANs are 3002 and 3003)
- If the physical NIC in your configuration carries the VLAN for VM, that VLAN is also present in the allowed VLAN list.

```
switch# show port-profile name alluplink
port-profile alluplink
  type: Ethernet
  description:
  status: enabled
  max-ports: 512
  min-ports: 1
  inherit:
  config attributes:
    switchport mode trunk
    switchport trunk allowed vlan 1,80,3002,610,620,630-650
  evaluated config attributes:
    switchport mode trunk
    switchport trunk allowed vlan 1,80,3002,3003,610,620,630-650
  no shutdown
  assigned interfaces:
    Ethernet2/2
  port-group:
  system vlans: none
  capability l3control: no
  capability iscsi-multipath: no
  capability vxlan: no
  capability l3-vn-service: no
  port-profile role: none
  port-binding: static
```

- Step 2** Verify that the Ethernet interface is up by entering the **ifconfig -a** command inside the VM. If not, consider deleting that NIC from the VM, and adding another NIC.
- Step 3** Using any sniffer tool, verify that ARP requests and responses are received on the VM interface.
- Step 4** On the upstream switch, look for the association between the IP and MAC address by entering these commands:
- **debug arp**
 - **show arp**

This example shows how to debug the Address Resolution Protocol (ARP):

```
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#
```

This example shows how to display ARP:

```
switch# show arp
Protocol  Address          Age (min)  Hardware Addr  Type   Interface
Internet  10.78.1.72       -          001a.6464.2008 ARPA   Vlan140
Internet  7.114.1.100     -          0011.bcac.6c00 ARPA   Vlan410
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
```

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
show mac address-table	Displays the MAC address table to verify all MAC addresses on all VEMs controlled by the VSM. See Example 8-1 on page 8-7 .
show mac address-table module <i>module-number</i>	Displays all the MAC addresses on the specified VEM.
show mac address-table static <i>HHHH.WWWW.HHHH</i>	Displays the MAC address table static entries. See Example 8-2 on page 8-8 .

Command	Purpose
show mac address-table address <i>HHHH.WWWW.HHHH</i>	Displays the interface on which the MAC address specified is learned or configured. <ul style="list-style-type: none"> For dynamic MAC addresses, if the same MAC address appears on multiple interfaces, then each of them is displayed separately. For static MAC addresses, if the same MAC address appears on multiple interfaces, then only the entry on the configured interface is displayed.
show mac address-table static inc veth	Displays the static MAC address of vEthernet interfaces in case a VEM physical port learns a dynamic MAC address and the packet source is in another VEM on the same VSM. See Example 8-3 on page 8-8 .
show running-config vlan <i>vlan-id</i>	Displays VLAN information in the running configuration.
show vlan [all-ports brief id <i>vlan-id</i> name <i>name</i> dot1q tag native]	Displays VLAN information as specified. See Example 8-4 on page 8-8 .
show vlan summary	Displays a summary of VLAN information.
show interface brief	Displays a table of interface states. See Example 8-5 on page 8-9 .
module vem <i>module-number</i> execute vemcmd show port	On the VEM, displays the port state on a particular VEM. This command can only be used from the VEM. See Example 8-6 on page 8-9 .
module vem <i>module-number</i> execute vemcmd show bd	For the specified VEM, displays its VLANs and their port lists. See Example 8-7 on page 8-9 .
module vem <i>module-number</i> execute vemcmd show trunk	For the specified VEM, displays the VLAN state on a trunk port. <ul style="list-style-type: none"> If a VLAN is forwarding (active) on a port, its CBL state should be 1. If a VLAN is blocked, its CBL state is 0. See Example 8-8 on page 8-10 .
module vem <i>module-number</i> execute vemcmd show l2 <i>vlan-id</i>	For the specified VEM, displays the VLAN forwarding table for a specified VLAN. See Example 8-9 on page 8-10 .
show interface <i>interface_id</i> mac-address	Displays the MAC addresses and the burn-in MAC address for an interface.

Example 8-1 show mac address-table command**Note**

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

**Tip**

Module indicates the VEM on which this MAC address is seen.

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

```
switch# show mac address-table
```

```

VLAN      MAC Address      Type   Age      Port
-----+-----+-----+-----+-----+-----
1         0002.3d20.2403   static 0        N1KV Internal Port      4
1         0002.3d30.2403   static 0        N1KV Internal Port      4
1         0002.3d40.2403   static 0        N1KV Internal Port      4
1         0002.3d60.2400   static 0        N1KV Internal Port      4
1         0002.3d80.2403   static 0        N1KV Internal Port      4
1         0000.0c07.accd   dynamic 1        Po3                      4
1         0050.56be.533f   dynamic 136     Po3                      4
1         4403.a74a.8422   dynamic 77      Po3                      4
1         4403.a74a.d586   dynamic 47      Po3                      4
1         5254.000a.ce25   dynamic 5        Po3                      4
1         5254.003e.8614   dynamic 5        Po3                      4
1         5254.0040.9ad6   dynamic 0        Po3                      4

```

switch#

Example 8-2 show mac address-table address command



Tip This command shows all interfaces on which a MAC address is learned dynamically. In this example, the same MAC address appears on Eth3/3 and Eth4/3.

```

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

```

Example 8-3 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 8-4 show vlan command



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

```

switch# show vlan
VLAN Name      Status      Ports
-----+-----+-----+-----+-----+-----
1      default     active     Po1, Po3, Po4, Eth4/1, Eth4/2
Eth5/1, Eth5/2, Eth6/1, Eth6/2
40     VLAN0040   active     Po1, Po3, Po4, Veth6, Veth7
Veth9, Eth4/1, Eth4/2, Eth5/1
Eth5/2, Eth6/1, Eth6/2

VLAN Type  Vlan-mode
----+-----+-----
1     enet    CE
40    enet    CE

```


Example 8-5 show interface brief command

```
switch# show interface brief
```

```
-----
Port      VRF      Status IP Address      Speed  MTU
-----
mgmt0    --      up    172.27.0.36      1000  1500
-----
Ethernet  VLAN    Type Mode    Status Reason      Speed  Port
Interface                                Speed  Ch #
-----
Eth4/1    1       eth trunk up    none      unknown  3
Eth4/2    1       eth trunk up    none      unknown  3
Eth5/1    1       eth trunk up    none      unknown  1
Eth5/2    1       eth trunk up    none      unknown  1
--More--2014 Sep 12 07:56:34 vsm-p last message repeated 6 times
```

Example 8-6 module vem module-number execute vemcmd show port command

Tip Look for the state of the port.

```
switch# module vem 3 execute vemcmd show port
vsm-p(vem-attach)# vemcmd show port
  LTL   VSM Port  Admin Link  State  PC-LTL  SGID          Vem Port  Type  ORG
svcpath Owner
  18   Eth4/1    UP   UP   FWD   1040    0           eth1     0
0
  19   Eth4/2    UP   UP   FWD   1040    1           eth0     0
0
  50   Veth6       UP   UP   FWD    0      0   cn1-vtep1-ovs  VXLAN  0
0
 1040   Po3         UP   UP   FWD    0          0
0
* 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 8-7 module vem 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 vem 5 execute vemcmd show bd

BD 1, vdc 1, vlan 1, swbd 1, 4 ports, ""

Portlist:
  12  _l24
  18  eth1
  19  eth0
 1040

BD 2, vdc 1, vlan 3972, swbd 3972, 0 ports, ""
```

```

Portlist:
BD 3, vdc 1, vlan 3970, swbd 3970, 0 ports, ""

Portlist:
BD 4, vdc 1, vlan 3968, swbd 3968, 3 ports, ""

Portlist:
   1 inband
   5 inband port security
  11 _l23

BD 5, vdc 1, vlan 3971, swbd 3971, 1 ports, ""

Portlist:
   15 _l27

BD 6, vdc 1, vlan 40, swbd 40, 4 ports, ""

Portlist:
   18 eth1
   19 eth0
   50 cn1-vtep1-ovs
  1040

```

Example 8-8 *module vem 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 vem 5 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(40) 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(40) cbl 1,
Trunk port 18 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(40) cbl 1,
Trunk port 19 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(40) cbl 1,
Trunk port 1040 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(40) cbl 1,

```

Example 8-9 *module vem module-number execute vemcmd show l2 command*

```

switch# configure terminal
n1000v(config)# module vem 3 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
switch#

```