



CHAPTER 9

Layer 2 Switching

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

This chapter includes the following sections:

- [Information About Layer 2 Ethernet Switching, page 9-1](#)
- [Port Model, page 9-1](#)
- [Layer 2 Switching Problems, page 9-4](#)
- [Verifying Layer 2 Switching, page 9-7](#)

Information About Layer 2 Ethernet Switching

Nexus1000V 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.

Port Model

This section describes the following port perspectives:


- [Viewing Ports from the VEM, page 9-2](#)
- [Viewing Ports from the VSM, page 9-3.](#)

[Send document comments to nexus1k-docfeedback@cisco.com.](mailto:nexus1k-docfeedback@cisco.com)

Viewing Ports from the VEM

The Nexus1000V differentiates between virtual and physical ports on each of the VEMs. [Figure 9-1](#) shows how ports on the Nexus1000V switch are bound to physical and virtual VMware ports within a VEM.

Figure 9-1 VEM View of Ports



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

- **Virtual NICs:** There are three types of Virtual NICs in VMware. The virtual NIC (vnic) is part of the VM, and represents the physical port of the host which is plugged into the switch. The virtual kernel NIC (vmknic) is used by the hypervisor for management, VMotion, iSCSI, NFS and other network access needed by the kernel. This interface would carry 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 of these types maps to a veth port within Nexus1000V.
- **Virtual Ethernet Ports (VEth):** A VEth port is a port on the Cisco Nexus 1000V Distributed Virtual Switch. 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.
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 are addressable by the module-port number method.

Send document comments to nexus1k-docfeedback@cisco.com.

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


- 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.
- Each uplink port on the host represents a physical interface. It acts a lot like an lvech port, but because physical ports do not move between hosts, the mapping is 1:1 between an uplink port and a vmnic.
- Each physical port added to Nexus1000V switch appears as a physical Ethernet port, just as it would on a hardware-based switch.

The uplink port concept is handled entirely by VMware, and is used to associate port configuration with vmnics. There is no fixed relationship between the uplink # and vmnic #, and these can be different on different hosts, and can change throughout the life of the host. On the VSM, the Ethernet interface number, such as ethernet 2/4, is derived from the vmnic number, not the uplink number.

Viewing Ports from the VSM

Figure 9-2 shows the VSM view ports.

Figure 9-2 VSM View of Ports



277344

[Send document comments to nexus1k-docfeedback@cisco.com.](mailto:nexus1k-docfeedback@cisco.com)

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 Virtual Machine on the ESX Host.
 - VMKNICs of the ESX Host
 - VSWIFs of an ESX COS Host.
- Eths (Physical Ethernet Interfaces) – correspond to the Physical NICs on the ESX Host.
- Po (Port Channel Interfaces) – The physical NICs of an ESX Host 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 Layer 2 Switching Configuration Guide, Release 4.0(4)SV1(3)*.

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 VEM Ports, page 9-4](#)
- [Verifying a Connection Between VEMs, page 9-5](#)
- [Isolating Traffic Interruptions, page 9-6](#)
- [Verifying Layer 2 Switching, page 9-7](#)

Verifying a Connection Between VEM Ports

To verify a connection between two veth ports on a VEM, follow these steps:

-
- Step 1** On the VSM, enter the **show vlan** command to view the state of the VLANs associated with the port. If the VLAN associated with a port is not active, then the port may 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 a **show interface brief** command.
 - Step 3** Enter the **module vem module-number execute vemcmd show port** command to 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.

The key things to look for in the output are:

- State of the port.
- CBL.
- Mode.
- Attached device name.
- The LTL of the port you are trying to troubleshoot. It will help you identify the interface quickly in other VEM commands where the interface name is not displayed.

Send document comments to nexus1k-docfeedback@cisco.com.

- Make sure the state of the port is up. If not, verify the configuration of the port on the VSM.
- Step 4** To view the VLANs and their port lists on a particular VEM, use the **module vem module-number execute vemcmd show bd** command:

```
n1000v# module vem 5 execute vemcmd show bd
```

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

Verifying a Connection Between VEMs

To verify a connection between veth ports on two separate VEMs, follow these steps:

- Step 1** Issue the **show vlan** command to check if the VLAN associated with the port is created on the VSM.
- Step 2** Issue the **show interface brief** command to check if the ports are up in the VSM.
- Step 3** On the VEM, issue the **module vem 3 execute vemcmd show port** command to check if the CBL state of the two ports is set to the value of 4 for forwarding.
- Step 4** On the VEM, issue the **module vem 3 execute vemcmd show bd** command to check if the two veth ports are listed in the flood list of the VLAN to which they are trying to communicate.
- 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 out the port on the upstream switch to which the pnuc (that is supposed to be carrying the VLAN) on the VEM is connected to.

```
n1000v# 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 you are looking for.

```
n1000v#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
```

Send document comments to nexus1k-docfeedback@cisco.com.

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

Use the following steps to isolate the cause for no traffic passing across VMs on different VEMs.

- Step 1** In 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 are 3002 and 3003)
 - If the physical NIC in your configuration carries the VLAN for VM, then that VLAN is also present in the allowed VLAN list.

```
n1000v#show port-profile name alluplink
port-profile alluplink
  description:
  status: enabled
  system vlans: 3002,3003
  port-group: alluplink
  config attributes:
    switchport mode trunk
    switchport trunk allowed vlan 1,80,3002,610,620,630-650
    no shutdown
  evaluated config attributes:
    switchport mode trunk
    switchport trunk allowed vlan 1,80,3002,3003,610,620,630-650
    no shutdown
  assigned interfaces:
    Ethernet2/2
```

- Step 2** Inside the VM, use the following command to verify that the Ethernet interface is up.

ifconfig -a

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, use the following commands to look for the association between the IP and MAC address:

debug arp

show arp

Example:

```
n1000v_CAT6K# 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
n1000v_CAT6K#
```

Example:

```
n1000v_CAT6K# sh arp
```

Send document comments to nexus1k-docfeedback@cisco.com.

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 5 You have completed this procedure.

Verifying Layer 2 Switching

Use the following commands to display and verify 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 9-1 on page 9-8
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 9-2 on page 9-9
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 MACs, if the same MAC appears on multiple interfaces, then each of them is displayed separately. For static MACs, if the same MAC appears on multiple interfaces, then only the entry on the configured interface is displayed.
show running-config vlan <vlan-id>	Displays VLAN information in the running configuration.
show vlan [all-ports brief id <vlan-id> name <name> dot1q tag native]	Displays VLAN information as specified. See Example 9-3 on page 9-9 .
show vlan summary	Displays a summary of VLAN information.
show interface brief	Displays a table of interface states. See Example 9-4 on page 9-10 .

Send document comments to nexus1k-docfeedback@cisco.com.

Command	Purpose
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 9-5 on page 9-10 .
module vem <i>module-number</i> execute vemcmd show bd command	For the specified VEM, displays its VLANs and their port lists. See Example 9-6 on page 9-11 .
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 active on a port, then its CBL state should be 4. • If a VLAN is blocked, then its CBL state is 1. See Example 9-7 on page 9-11 .
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 9-8 on page 9-11 .
show interface <i>interface_id</i> mac	Displays the MAC addresses and the burn-in MAC address for an interface.

Example 9-1 show mac address-table Command



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



Tip VEM indicates on which VEM this MAC is seen.

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

```
n1000v# show mac address-table
Legend:
      * - primary entry, G - Gateway MAC, (R) - Routed MAC
      age - seconds since last seen
      VEM      VLAN      MAC Address      Type      age      Ports
-----+-----+-----+-----+-----+-----+-----
*    3        1        0002.3d22.e300   static    -        N1KV Internal Port
*    3        1        0002.3d22.e302   static    -        N1KV Internal Port
*    4        1        0002.3d22.e303   static    -        N1KV Internal Port
*    3        1        0002.3d32.e300   static    -        N1KV Internal Port
*    3        1        0002.3d32.e302   static    -        N1KV Internal Port
*    4        1        0002.3d32.e303   static    -        N1KV Internal Port
*    3        1        0002.3d62.e300   static    -        N1KV Internal Port
*    3        1        0002.3d62.e302   static    -        N1KV Internal Port
*    4        1        0002.3d62.e303   static    -        N1KV Internal Port
      4        1        0023.7d34.f4e2   dynamic   23        Eth4/2
      3        115       0002.3d42.e302   dynamic    0        N1KV Internal Port
      4        115       0002.3d42.e303   dynamic    0        N1KV Internal Port
```


Send document comments to nexus1k-docfeedback@cisco.com.

```

4      115      0050.56bb.49d9   dynamic  0      Eth4/2
3      115      0050.56bb.49d9   dynamic  0      Eth3/4
3      116      0002.3d22.e302   dynamic  1      N1KV Internal Port
4      116      0002.3d22.e302   dynamic  1      Eth4/2
4      116      0002.3d22.e303   dynamic  1      N1KV Internal Port
3      116      0002.3d22.e303   dynamic  1      Eth3/4

```

Example 9-2 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 Eth4/2 and Eth3/4.

```

n1000v# show mac address-table address 0050.56bb.49d9
Legend:
      * - primary entry, G - Gateway MAC, (R) - Routed MAC
      age - seconds since last seen

```

VEM	VLAN	MAC Address	Type	age	Ports
4	115	0050.56bb.49d9	dynamic	0	Eth4/2
3	115	0050.56bb.49d9	dynamic	0	Eth3/4

Example 9-3 show vlan Command



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

```

n1000v# show vlan

```

VLAN	Name	Status	Ports
1	default	active	Eth3/3, Eth3/4, Eth4/2, Eth4/3
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	

Send document comments to nexus1k-docfeedback@cisco.com.

```

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 9-4 show interface brief Command

```

n1000v# show int brief
-----
Port      VRF      Status IP Address      Speed  MTU
-----
mgmt0    --      up      172.23.232.143  1000  1500
-----

Ethernet  VLAN  Type Mode  Status Reason      Speed  Port
Interface
-----
Eth3/4    1     eth trunk up     none      1000 (D) --
Eth4/2    1     eth trunk up     none      1000 (D) --
Eth4/3    1     eth trunk up     none      1000 (D) --

```

Example 9-5 module vem module-number execute vemcmd show port Command



Tip Look for the state of the port.

```

~ # module vem 3 execute vemcmd show port
LTL   IfIndex  Vlan   Bndl  SG_ID  Pinned_SGID  Type  Admin  State  CBL  Mode  Name
  8     0     3969   0     2     2     VIRT  UP     UP     4   Access  120
  9     0     3969   0     2     2     VIRT  UP     UP     4   Access  121
 10     0     115    0     2     0     VIRT  UP     UP     4   Access  122
 11     0     3968   0     2     2     VIRT  UP     UP     4   Access  123
 12     0     116    0     2     0     VIRT  UP     UP     4   Access  124
 13     0     1      0     2     2     VIRT  UP     UP     0   Access  125
 14     0     3967   0     2     2     VIRT  UP     UP     4   Access  126
 16   1a030100    1 T     0     0     2     PHYS  UP     UP     4   Trunk
vmnic1
 17   1a030200    1 T     0     2     2     PHYS  UP     UP     4   Trunk
vmnic2

```

Send document comments to nexus1k-docfeedback@cisco.com.

Example 9-6 *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.

```

~ # module vem 5 execute vemcmd show bd
Number of valid BDS: 8
BD 1, vdc 1, vlan 1, 2 ports
Portlist:
16 vmmic1
17 vmmic2
BD 100, vdc 1, vlan 100, 0 ports
Portlist:
BD 110, vdc 1, vlan 110, 1 ports
Portlist:
16 vmmic1
BD 111, vdc 1, vlan 111, 1 ports
Portlist:
16 vmmic1
BD 112, vdc 1, vlan 112, 1 ports
Portlist:
16 vmmic1
BD 113, vdc 1, vlan 113, 1 ports
Portlist:
16 vmmic1
BD 114, vdc 1, vlan 114, 1 ports
Portlist:
16 vmmic1
BD 115, vdc 1, vlan 115, 2 ports
Portlist:
10 l22
16 vmmic1

```

Example 9-7 *module vem module-number execute vemcmd show trunk Command*



Tip If a VLAN is active on a port, then its CBL state should be 4.
If a VLAN is blocked, then its CBL state is 1.

```

~ # module vem 5 execute vemcmd show trunk
Trunk port 16 native_vlan 1 CBL 4
vlan(1) cbl 4, vlan(110) cbl 4, vlan(111) cbl 4, vlan(112) cbl 4, vlan(113) cbl 4,
vlan(114) cbl 4, vlan(115) cbl 4, vlan(116) cbl 4, vlan(117) cbl 4, vlan(118) cbl 4,
vlan(119) cbl 4,
Trunk port 17 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(117) cbl 4,
~ #

```

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

```

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

```

Send document comments to nexus1k-docfeedback@cisco.com.