



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.

Port Model

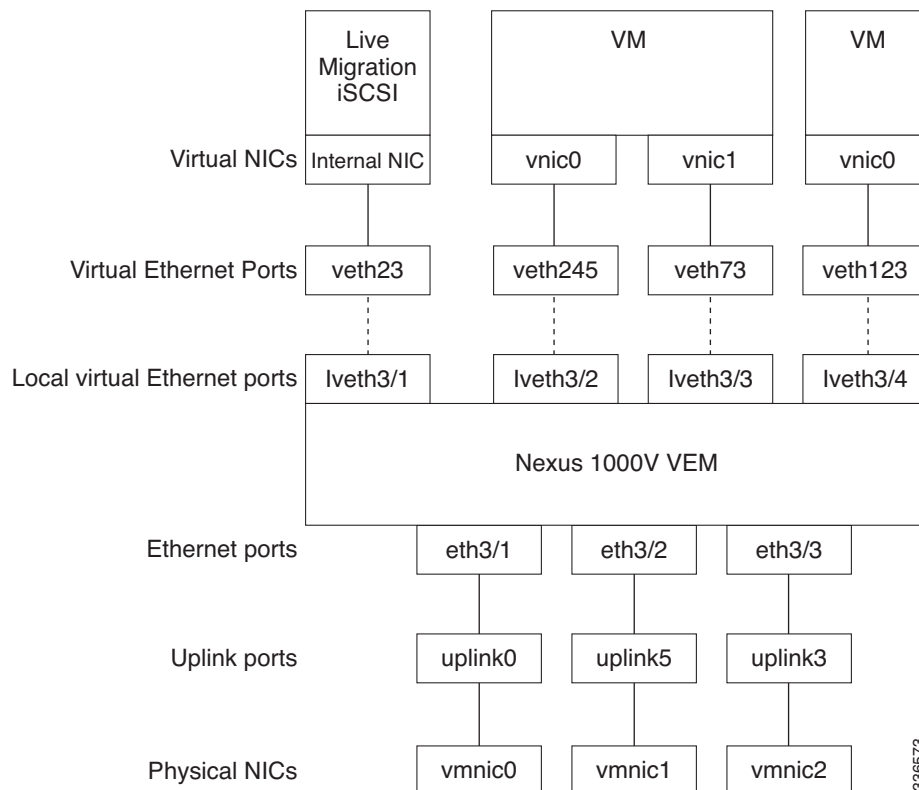
This section describes the following port perspectives:

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

Viewing Ports from the VEM

The Cisco Nexus 1000V differentiates between virtual and physical ports on each of the VEMs. [Figure 10-1](#) shows how ports on the Cisco Nexus 1000V switch are bound to physical and virtual Microsoft Hyper-V ports within a VEM.

Figure 10-1 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 called a vmnic. The vmnic number is allocated during Microsoft Hyper-V 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. The port acts like a local vEth 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 that is added to the Cisco Nexus 1000V switch appears as a physical Ethernet port, just as it would on a hardware-based switch.

The uplink port concept is handled entirely by the hypervisor and is used to associate port configuration with vmnics. There is no fixed relationship between the uplink number and vmnic number, and the uplink and the vmnic numbers 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 10-2 shows the VSM view of the ports.

Figure 10-2 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 Virtual Machine 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 Microsoft Hyper-V Layer 2 Switching Configuration Guide, Release 5.2(1)SM1(5.1)*.

Problems with Layer 2 Switching

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

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.

```
n1000V# 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.

```
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 that 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
```

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.

```
n1000v# 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):

```
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#
```

This example shows how to display ARP:

```
n1000v_CAT6K# 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

```

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 10-1 on page 10-8 .
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 10-2 on page 10-9 .
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 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 10-3 on page 10-9 .
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 10-4 on page 10-9 .
show vlan summary	Displays a summary of VLAN information.
show interface brief	Displays a table of interface states. See Example 10-5 on page 10-10 .

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 10-6 on page 10-10 .
module vem <i>module-number</i> execute vemcmd show bd	For the specified VEM, displays its VLANs and their port lists. See Example 10-7 on page 10-11 .
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 forwarding (active) on a port, its CBL state should be 1. • If a VLAN is blocked, its CBL state is 0. See Example 10-8 on page 10-11 .
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 10-9 on page 10-12 .
show interface <i>interface_id</i> mac-address	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 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.

```
n1000v# show mac address-table
VLAN      MAC Address      Type    Age    Port                               Mod
-----+-----+-----+-----+-----+-----+-----
1         0002.3d11.5502   static  0      N1KV Internal Port                3
1         0002.3d21.5500   static  0      N1KV Internal Port                3
1         0002.3d21.5502   static  0      N1KV Internal Port                3
1         0002.3d31.5502   static  0      N1KV Internal Port                3
1         0002.3d41.5502   static  0      N1KV Internal Port                3
1         0002.3d61.5500   static  0      N1KV Internal Port                3
1         0002.3d61.5502   static  0      N1KV Internal Port                3
1         0002.3d81.5502   static  0      N1KV Internal Port                3
3         12ab.47dd.ff89   static  0      Eth3/3                             3
342       0002.3d41.5502   static  0      N1KV Internal Port                3
342       0050.568d.5a3f   dynamic 0      Eth3/3                             3
343       0002.3d21.5502   static  0      N1KV Internal Port                3
343       0050.568d.2aa0   dynamic 9      Eth3/3                             3
Total MAC Addresses: 13
n1000v#
```


Example 10-2 *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.

```
n1000v# 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
n1000v#
```

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

```
n1000v# show mac address-table static | inc veth
460      0050.5678.ed16  static 0      Veth2     3
460      0050.567b.1864  static 0      Veth1     4
n1000v#
```

Example 10-4 *show vlan command*

Tip This command shows the state of each VLAN that is 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
815  VLAN0815                active
816  VLAN0816                active
```

```

817 VLAN0817          active
818 VLAN0818          active
819 VLAN0819          active
820 VLAN0820          active

```

```

VLAN Type Vlan-mode
-----

```

```

Remote SPAN VLANs
-----

```

```

Primary Secondary Type          Ports
-----

```

Example 10-5 show interface brief command

```
n1000v# show interface 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                               Speed  Ch #
-----
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) --
-----

Port-channel VLAN  Type Mode  Status Reason          Speed  Protocol
Interface
-----
Po1        1     eth trunk up    none           a-1000(D) none
Po2        1     eth pvlan up    none           a-10G(D) none
-----

Vethernet VLAN  Type Mode  Status Reason          Speed
-----
Veth1      262  virt access up    none           auto
-----

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

```

Example 10-6 module vem module-number execute vemcmd show port command



Tip Look for the state of the port.

```

n1000v# 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    1  Access 120
9    0    3969  0     2      2    VIRT  UP    UP    1  Access 121
10   0    115   0     2      0    VIRT  UP    UP    1  Access 122
11   0    3968  0     2      2    VIRT  UP    UP    1  Access 123

```

12	0	116	0	2	0	VIRT	UP	UP	1	Access	l24	
13	0	1	0	2	2	VIRT	UP	UP	0	Access	l25	
14	0	3967	0	2	2	VIRT	UP	UP	1	Access	l26	
16	1a030100	1	T	0	0	2	PHYS	UP	UP	1	Trunk	vmnic1
17	1a030200	1	T	0	2	2	PHYS	UP	UP	1	Trunk	vmnic2

Example 10-7 *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.

```
n1000v# module vem 5 execute vemcmd show bd
Number of valid BDS: 8
BD 1, vdc 1, vlan 1, 2 ports
Portlist:
16 vmnic1
17 vmnic2
BD 100, vdc 1, vlan 100, 0 ports
Portlist:
BD 110, vdc 1, vlan 110, 1 ports
Portlist:
16 vmnic1
BD 111, vdc 1, vlan 111, 1 ports
Portlist:
16 vmnic1
BD 112, vdc 1, vlan 112, 1 ports
Portlist:
16 vmnic1
BD 113, vdc 1, vlan 113, 1 ports
Portlist:
16 vmnic1
BD 114, vdc 1, vlan 114, 1 ports
Portlist:
16 vmnic1
BD 115, vdc 1, vlan 115, 2 ports
Portlist:
10 l22
16 vmnic1
```

Example 10-8 *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.

```
n1000v# module vem 5 execute vemcmd show trunk
Trunk port 16 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(110) cbl 1, vlan(111) cbl 1, vlan(112) cbl 1, vlan(113) cbl 1,
vlan(114) cbl 1, vlan(115) cbl 1, vlan(116) cbl 1, vlan(117) cbl 1, vlan(118) cbl 1,
vlan(119) cbl 1,
Trunk port 17 native_vlan 1 CBL 0
vlan(1) cbl 1, vlan(117) cbl 1,
n1000v#
```

Example 10-9 *module vem module-number execute vemcmd show l2 command*

```
n1000v# 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
n1000v#
```

Troubleshooting Microsoft NLB Unicast Mode

Microsoft Network Load Balancing (MS-NLB) is a clustering technology offered by Microsoft as part of the Windows server operating systems. Clustering enables a group of independent servers to be managed as a single system for higher availability, easier manageability, and greater scalability.

For more information about MS-NLB, see the following URL:

<http://technet.microsoft.com/en-us/library/bb742455.aspx>



Note

Access to third-party websites identified in this document is provided solely as a courtesy to customers and others. Cisco Systems, Inc. and its affiliates are not in any way responsible or liable for the functioning of any third-party website, or the download, performance, quality, functioning or support of any software program or other item accessed through the website, or any damages, repairs, corrections or costs arising out of any use of the website or any software program or other item accessed through the website. Cisco's End User License Agreement does not apply to the terms and conditions of use of a third-party website or any software program or other item accessed through the website.

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:

- Private VLAN (PVLAN) port
- Ports configured with unknown unicast flood blocking (UUFB)
- Ports configured with a switchport port-security mac-address sticky

Disabling Automatic Static MAC Learning on vEthernet Interfaces

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

In interface configuration mode, enter these commands:

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

In port profile configuration mode, enter these commands:

```
switch(config)# port-profile type vethernet ms-nlb
switch(config-port-prof)# no mac auto-static-learn
```

Checking the Status on a VSM

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

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

This example shows how to generate the VSM status in the interface configuration mode:

```
switch(config-if)# show running-config int veth1
interface Vethernet1
  inherit port-profile vm59
  description Fedora117, Network Adapter 2
  switchport port-security mac-address 001D.D8B7.1F81
  dvport uuid "ea 5c 3b 50 cd 00 9f 55-41 a3 2d 61 84 9e 0e c4"
```

This example shows how to generate the VSM status in the port profile configuration:

```
switch(config-if)# show running-config port-profile ms-nlb
port-profile type vethernet ms-nlb
  ip port access-group abhi-acl in
  ip port access-group abhi-acl out
  no shutdown
  guid a85154f4-b07a-4cc4-86ad-fac0246557fe
  publish port-profile
  max-ports 300
  state enabled
```

Checking the Status on a VEM

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

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

This example shows how to generate the VSM status:

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

This example shows how to 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
```

Configuring UUFb to Block Unwanted MS-NLB Traffic

When MS NLB VMs have more than one port on the same subnet, a request is flooded, which causes both ports to receive it. The server cannot manage this situation.

A workaround for this situation is to enable unknown unicast flood blocking (UUFb).

Enabling UUFb

This example shows how to enable UUFb. After you enter the commands in the example, press Ctrl-Z.

```
n1000v# configure terminal
n1000v (config)# uufb enable
n1000v (config)#
```

This configuration conceals the requests from the non-NLB ports and allows the system to function as expected.

Disabling UUFb for VMs That Use Dynamic MAC Addresses

Issues might occur for VMs that use dynamic MAC addresses. For ports that host these types of VMs, disable UUFb.

This example shows how to disable UUFb:

```
n1000v(config)# interface veth3
n1000v(config-if)# switchport uufb disable
n1000v(config-if)#
```