

# Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Configure](#)

[Network Topology](#)

[Tracing of MAC Addresses at Different Network Segments](#)

[Verify](#)

[Troubleshoot](#)

## Introduction

This document describes how to trace MAC addresses of a virtual machine (VM) and VMkernel (VMK) interface at these network levels:

- Cisco Nexus 5000 Series Switches
- Cisco Unified Computing System (UCS) 6248 Fabric Interconnect (FI)
- VMware ESXi host
- Cisco Nexus 1000V Switch

It is important to understand which uplink a VM or VMK interface uses for communication for both troubleshooting and design aspects.

## Prerequisites

### Requirements

Cisco recommends that you have knowledge of these topics:

- vPC feature in Cisco NX-OS
- Cisco Unified Computing System
- VMware ESXi
- Cisco Nexus 1000V Switch

### Components Used

The information in this document is based on these software and hardware versions:

- Cisco Nexus 5020 Switch Version 5.0(3)N2(2a)
- Cisco Unified Computing System Version 2.1(1d)
- Cisco Unified Computing System B200 M3 Blade Server with Cisco Virtual Interface Card (VIC) 1240 (Palo) CNAvSphere 5.1 (ESXi and vCenter)
- Cisco Nexus 1000V Switch Version 4.2(1)SV2(1.1a)

The information in this document was created from the devices in a specific lab environment. All of

the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

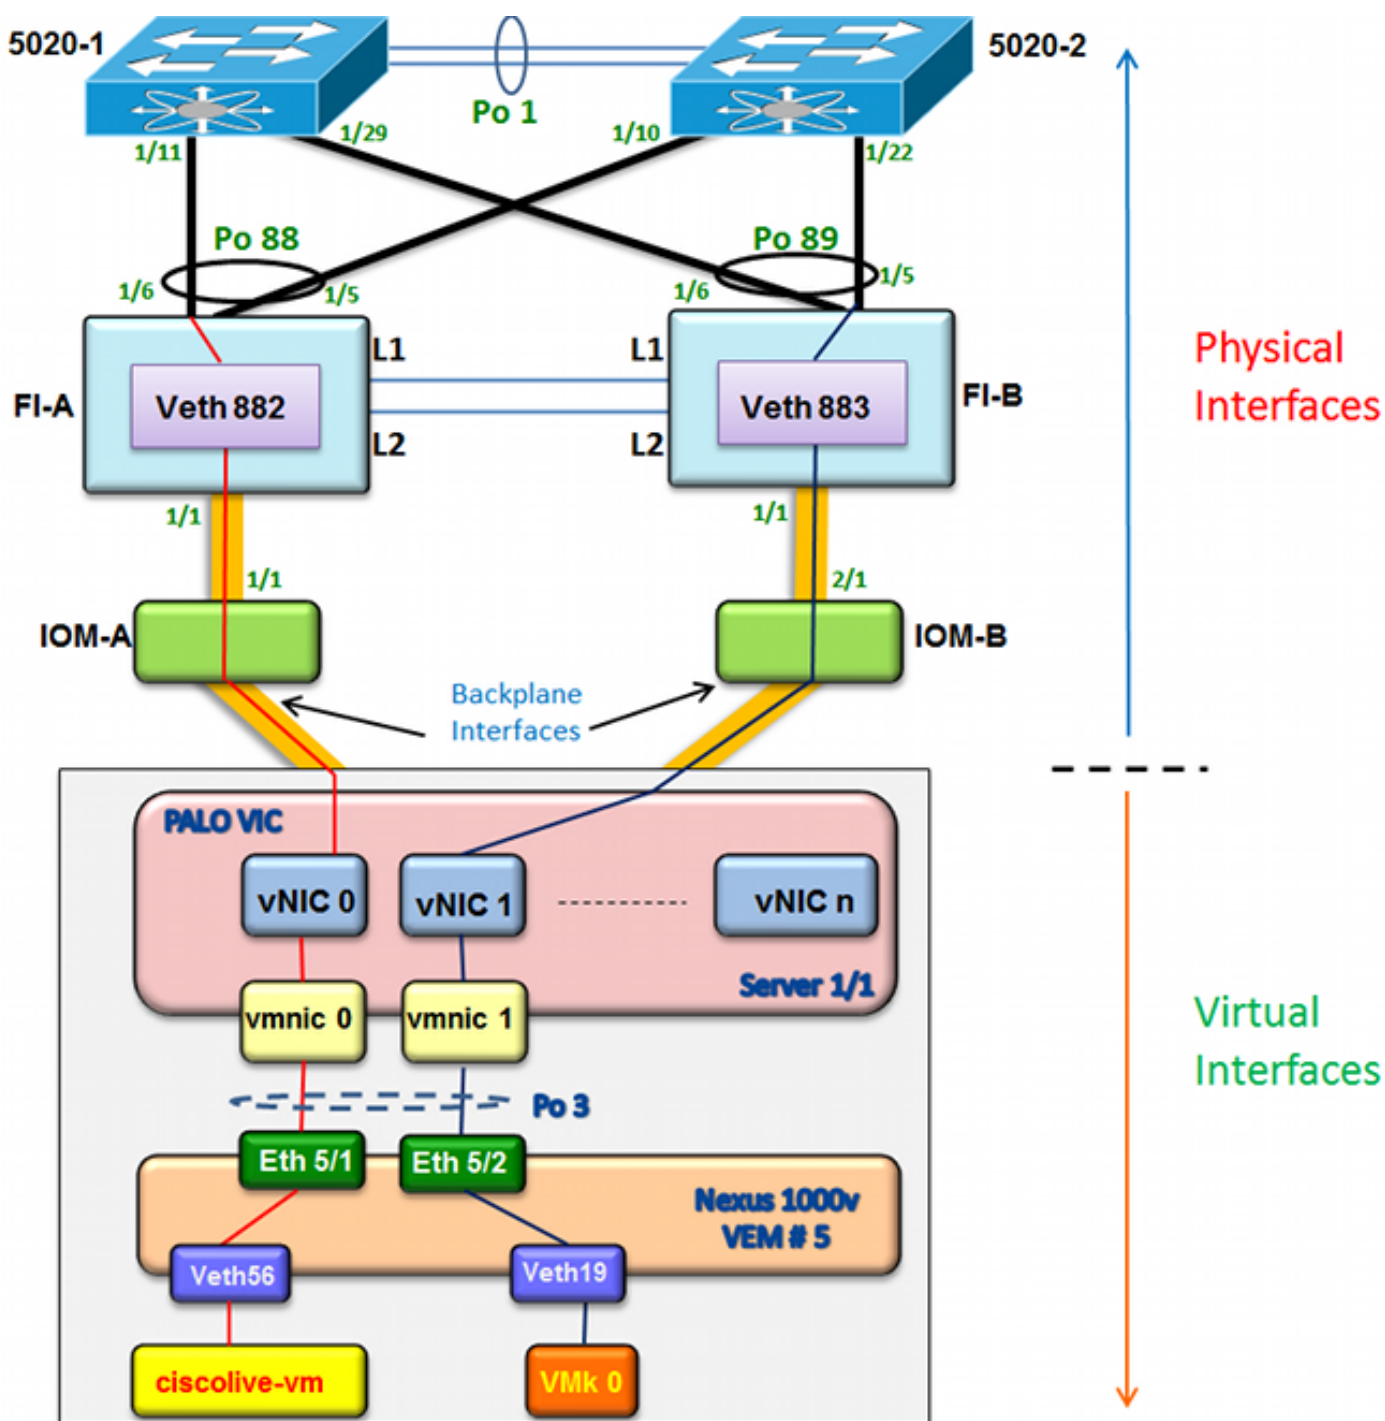
## Configure

### Network Topology

In this example setup, the VM and VMK interfaces are on same host (IP address 172.16.18.236) and the same VLAN 18 (subnet 172.16.18.0/24).

In the Nexus 1000V, the host is represented as Virtual Ethernet Module (VEM) # 5.

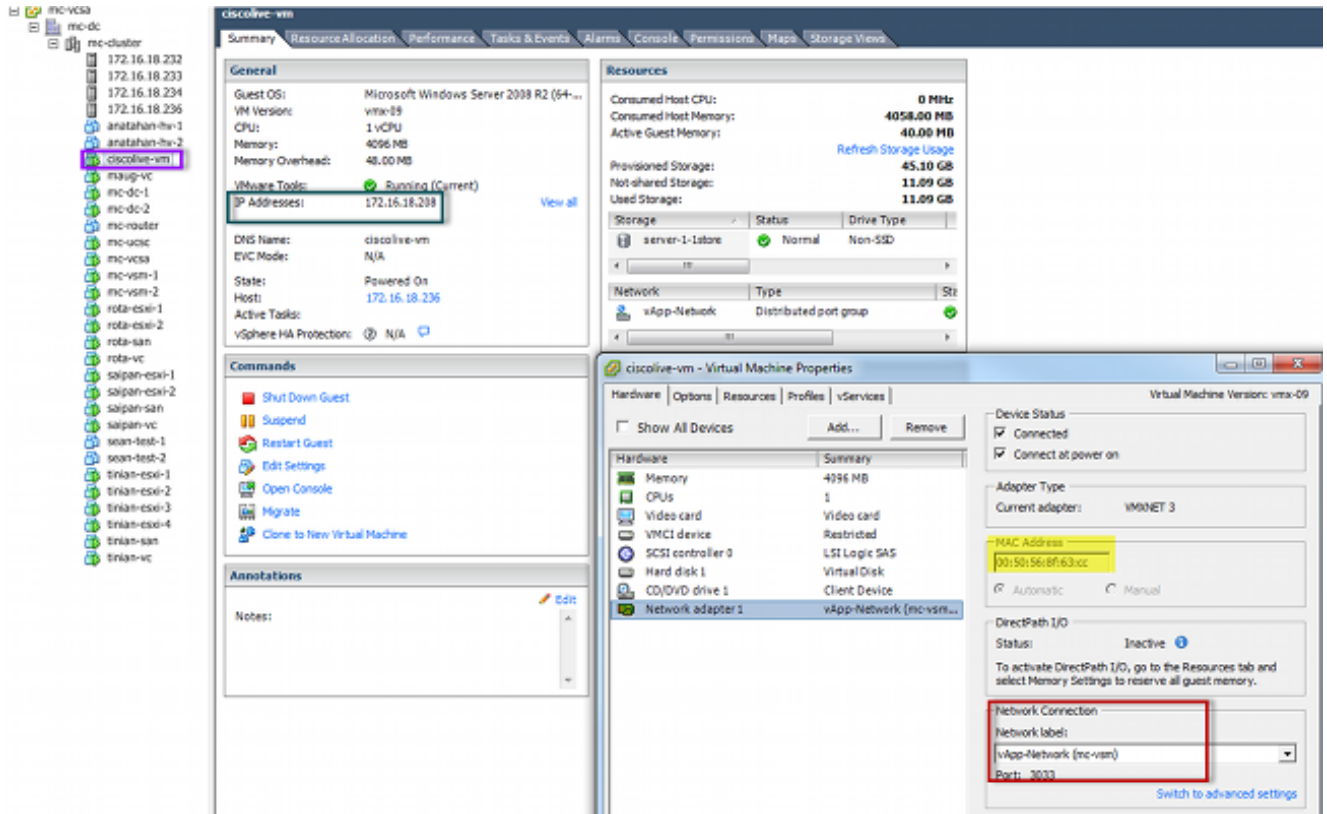
In the UCS, the host is installed on blade 1 in chassis 1.



# Tracing of MAC Addresses at Different Network Segments

This procedure describes how to trace MAC addresses at various network levels.

1. In the vCenter, find the MAC address of the VM you wish to trace. In this example, the MAC address of the VM (ciscolive-vm) is 0050:568f:63cc:




2. Enter the **esxcfg-vmknic -l** command on the ESXi shell in order to find the MAC address of VMK interface from the host. In this example, the VMK (vmk0) is the management interface and has a MAC address of 0050:56:67:8e:b9:

```
mc-vsm# show mac address-table | in 8eb9
18      0050.5667.8eb9      static  0      Veth19      5
18      0050.5667.8eb9      dynamic 0      Po4         6
mc-vsm# show mac address-table | in 63cc
18      0050.568f.63cc      dynamic 93     Po1         3
18      0050.568f.63cc      dynamic 93     Po2         4
18      0050.568f.63cc      static  0      Veth56      5
18      0050.568f.63cc      dynamic 93     Po4         6
mc-vsm#
```

3. Confirm that the MAC addresses of the VM (ciscolive-vm) and the VMK interface (vmk0) are being learned on the ESXi host (VEM) and the Nexus 1000V.

At the VEM level, enter the **vemcmd show I2 18** command in order to confirm that both MAC addresses are learned:



At the Nexus 1000V level, enter a **show mac address-table** command in order to confirm that both MAC addresses are learned on VLAN 18 on VEM # 5:

```
mc-vsm# show mac address-table | in 8eb9
18      0050.5667.8eb9    static 0          Veth19          5
18      0050.5667.8eb9    dynamic 0         Po4              6
mc-vsm# show mac address-table | in 63cc
18      0050.568f.63cc    dynamic 93         Po1              3
18      0050.568f.63cc    dynamic 93         Po2              4
18      0050.568f.63cc    static 0          Veth56          5
18      0050.568f.63cc    dynamic 93         Po4              6
mc-vsm#
```

Enter the **show port-channel summary** command for VEM # 5 in order to see the port-channel and member ports:

```
mc-vsm#
mc-vsm# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
       I - Individual    H - Hot-standby (LACP only)
       s - Suspended     r - Module-removed
       S - Switched      R - Routed
       U - Up (port-channel)

-----
Group  Port-      Type   Protocol  Member Ports
Channel
-----
1      Po1 (SU)   Eth    NONE      Eth3/1 (P)  Eth3/2 (P)  Eth3/9 (r)
                               Eth3/10 (r)
2      Po2 (SU)   Eth    NONE      Eth4/1 (P)  Eth4/2 (P)  Eth4/9 (P)
                               Eth4/10 (P)
3      Po3 (SU)   Eth    NONE      Eth5/1 (P)  Eth5/2 (P)  Eth5/9 (r)
                               Eth5/10 (r)
4      Po4 (SU)   Eth    NONE      Eth6/1 (P)  Eth6/2 (P)  Eth6/11 (P)
                               Eth6/12 (P)
```

4. Gather additional details from Nexus 1000V.

Enter the **show interface vethernet 56** command in order to see that Veth56 corresponds to the VM (ciscolive-vm):

```
mc-vsm# show interface vethernet 56
Vethernet56 is up
Port description is ciscolive-vm, Network Adapter 1
Hardware: Virtual, address: 0050.568f.63cc (bia 0050.568f.63cc)
Owner is VM "ciscolive-vm", adapter is Network Adapter 1
Active on module 5
VMware DVS port 3033
Port-Profile is vApp-Network
Port mode is access
5 minute input rate 80 bits/second, 0 packets/second
5 minute output rate 12552 bits/second, 8 packets/second
Rx
 23795 Input Packets 7293075158593488853 Unicast Packets
 203449390 Multicast Packets 4294967761 Broadcast Packets
 2333878 Bytes
Tx
 1350625 Output Packets 4768 Unicast Packets
 519692101807 Multicast Packets 4321524090 Broadcast Packets 1345857 Flood Packets
 254466737 Bytes
 0 Input Packet Drops 0 Output Packet Drops
```

Enter the **show interface vethernet 19** command in order to see that Veth19 corresponds to the VMK interface (vmk0) of the host:

```
mc-vsm# show interface vethernet 19
Vethernet19 is up
Port description is VMware VMkernel, vmk0
Hardware: Virtual, address: 0050.5667.8eb9 (bia 0050.5667.8eb9)
Owner is VMware VMkernel, adapter is vmk0
Active on module 5
VMware DVS port 2110
Port-Profile is 13
Port mode is access
5 minute input rate 12904 bits/second, 1 packets/second
5 minute output rate 13384 bits/second, 8 packets/second
Rx
 12200 Input Packets 7310589476873731518 Unicast Packets
 7310589476867241067 Multicast Packets 873444753044241742 Broadcast Packets
 16040625 Bytes
Tx
 65549 Output Packets 3731 Unicast Packets
 141938759046 Multicast Packets 137454132371 Broadcast Packets 59221 Flood Packets
 12416427 Bytes
 8227343645136678255 Input Packet Drops 210453427045 Output Packet Drops
```

5. Check the pinning of the traffic from the VM (ciscolive-vm) and the VMK interface

```
mc-vsm# module vem 5 execute vemcmd show port vsm
```

LTL	VSM Port	Admin	Link	State	PC-LTL	SGID	Vem Port	Type
6	Internal	DOWN	UP	FWD	0		vns	
8	Internal	UP	UP	FWD	0			
9	Internal	DOWN	DOWN	FWD	0			
10	Internal	DOWN	DOWN	FWD	0	0		
11	Internal	DOWN	DOWN	FWD	0	0		
12	Internal	DOWN	DOWN	FWD	0	0		
14	Internal	DOWN	DOWN	FWD	0			
15	Internal	DOWN	DOWN	FWD	0			
16	Internal	DOWN	DOWN	FWD	0			ar
17	Eth5/1	UP	UP	FWD	305	0	vmnic0	
18	Eth5/2	UP	UP	FWD	305	1	vmnic1	
49	Veth19	UP	UP	FWD	0	1	vmk0	
50	Veth23	UP	UP	FWD	0	1	tinian-san.eth0	
51	Veth38	UP	UP	F/B*	0	0	tinian-esxi-1.eth3	
52	Veth37	UP	UP	F/B*	0	0	tinian-esxi-1.eth2	
53	Veth22	UP	UP	F/B*	0	1	tinian-esxi-1.eth1	
54	Veth21	UP	UP	F/B*	0	0	tinian-esxi-1.eth0	
55	Veth36	UP	UP	F/B*	0	1	tinian-esxi-2.eth3	
56	Veth35	UP	UP	F/B*	0	0	tinian-esxi-2.eth2	
57	Veth25	UP	UP	F/B*	0	1	tinian-esxi-2.eth1	
58	Veth24	UP	UP	F/B*	0	0	tinian-esxi-2.eth0	
59	Veth43	UP	UP	F/B*	0	1	tinian-esxi-3.eth3	
60	Veth44	UP	UP	F/B*	0	0	tinian-esxi-3.eth2	
61	Veth45	UP	UP	F/B*	0	1	tinian-esxi-3.eth1	
62	Veth46	UP	UP	F/B*	0	0	tinian-esxi-3.eth0	
63	Veth47	UP	UP	F/B*	0	1	tinian-esxi-4.eth3	
64	Veth48	UP	UP	F/B*	0	0	tinian-esxi-4.eth2	
65	Veth49	UP	UP	F/B*	0	1	tinian-esxi-4.eth1	
66	Veth50	UP	UP	F/B*	0	0	tinian-esxi-4.eth0	
67	Veth26	UP	UP	FWD	0	1	tinian-vc.eth0	
68	Veth56	UP	UP	FWD	0	0	ciscolive-vm.eth0	
69	Veth31	UP	UP	FWD	0	1	maug-vc.eth0	
75	Veth59	UP	UP	FWD	0	0	mc-ucsc.eth0	
78	Veth72	UP	UP	FWD	0	1	mc-dc-2.eth0	
305	Po3	UP	UP	FWD				

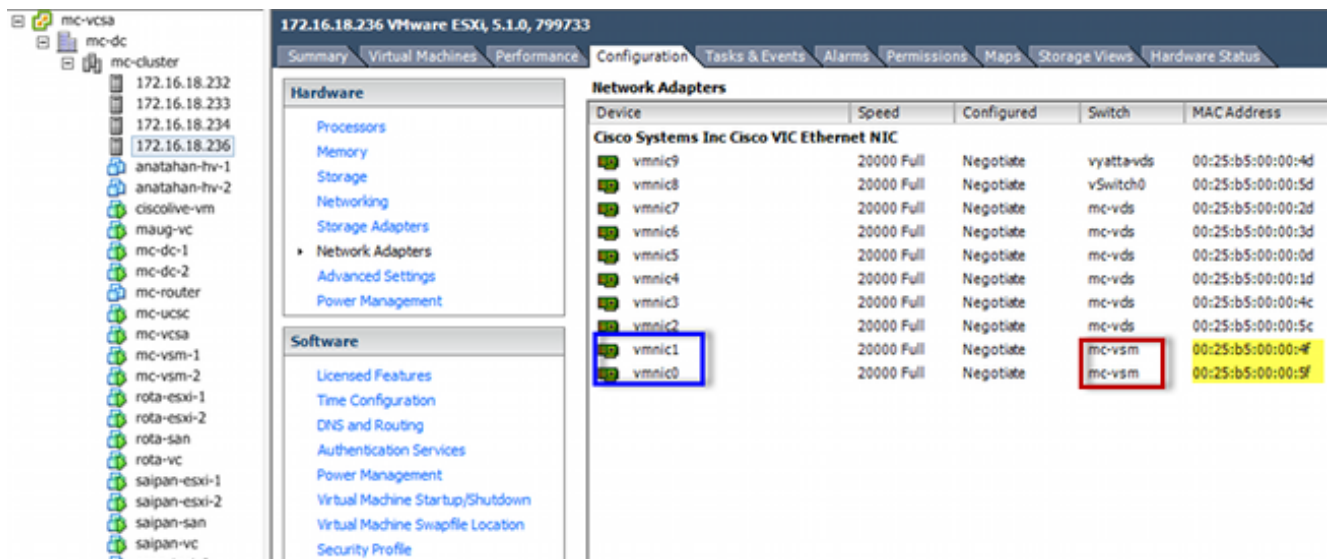
\* F/B: Port is BLOCKED on some of the vlans.  
 One or more vlans are either not created or  
 not in the list of allowed vlans for this port.  
 Please run "vemcmd show port vlans" to see the details.  
 mc-vsm#

This output shows the subscriber group ID (SGID) mapping for the VM (ciscolive-vm) and the VMK interface

- SGID 0 of the VM (ciscolive-vm) matches SGID 0 of vmnic0.
- SGID 1 of the VMK interface

6. Get the MAC addresses of the VMNICs from either the vCenter or the ESXi command-line interface (CLI).

In the vCenter, navigate to the **Configuration** tag:

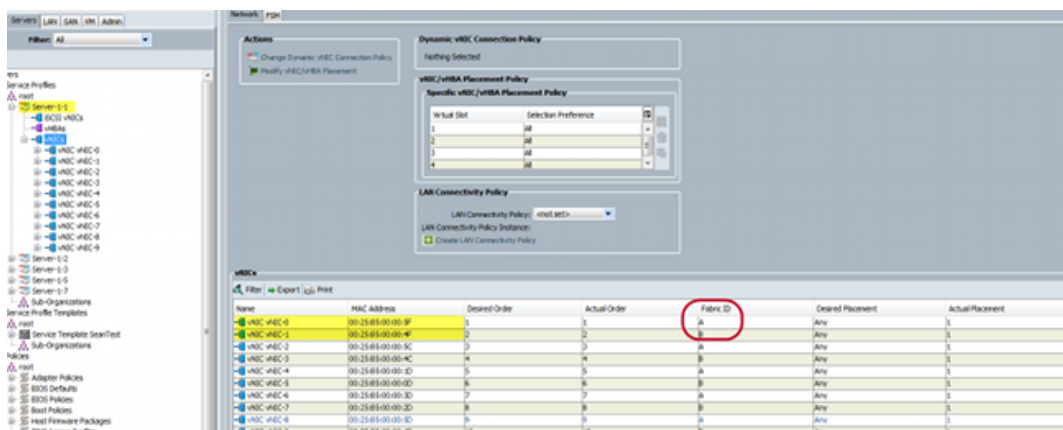


On the ESXi CLI, enter the **esxcfg-nics -l** command:

```

- # esxcfg-nics -l
Name PCI Driver Link Speed Duplex MAC Address MTU Description
vmnic0 0000:06:00.00 enic Up 20000Mbps Full 00:25:b5:00:00:5f 1500 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic1 0000:07:00.00 enic Up 20000Mbps Full 00:25:b5:00:00:4f 1500 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic2 0000:08:00.00 enic Up 20000Mbps Full 00:25:b5:00:00:5c 9000 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic3 0000:09:00.00 enic Up 20000Mbps Full 00:25:b5:00:00:4c 9000 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic4 0000:0a:00.00 enic Up 20000Mbps Full 00:25:b5:00:00:1d 9000 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic5 0000:0b:00.00 enic Up 20000Mbps Full 00:25:b5:00:00:0d 9000 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic6 0000:0c:00.00 enic Up 20000Mbps Full 00:25:b5:00:00:3d 9000 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic7 0000:0d:00.00 enic Up 20000Mbps Full 00:25:b5:00:00:2d 9000 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic8 0000:0e:00.00 enic Up 20000Mbps Full 00:25:b5:00:00:5d 9000 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic9 0000:0f:00.00 enic Up 20000Mbps Full 00:25:b5:00:00:4d 9000 Cisco Systems Inc Cisco VIC Ethernet NIC
  
```

7. In the UCS Manager (UCSM), find the virtual network interface controllers (vNICs) of the UCS that correspond to the VMNICs:



The primary FI for vNIC-0 is FI-A, while the primary FI for vNIC-1 is FI-B. You can now infer that traffic from the VM (ciscolive-vm) traverses through FI-A and that traffic from the VMK interface (vmk0) traverses through FI-B.

8. Confirm that the MAC address of the VM (ciscolive-vm) is learned on FI-A:

```
Mike-Cliff-Pod-16-A(nxos)# show mac address-table | in 63cc
* 18      0050.568f.63cc      dynamic  0          F    F    Veth882
Mike-Cliff-Pod-16-A(nxos)#
Mike-Cliff-Pod-16-A(nxos)# show int vethernet 882
Vethernet882 is up
  Bound Interface is port-channel1288
  Hardware: Virtual, address: 547f.eea2.5ac0 (bia 547f.eea2.5ac0)
  Description: server 1/1, VNIC vNIC-0
  Encapsulation ARPA
  Port mode is trunk
  EtherType is 0x8100
  Rx
    38196726 unicast packets  130708 multicast packets  99167 broadcast packets
    38426601 input packets  44470647026 bytes
    0 input packet drops
  Tx
    18711011 unicast packets  552876 multicast packets  10560283 broadcast packets
    29824170 output packets  9379742901 bytes
    0 flood packets
    0 output packet drops
```

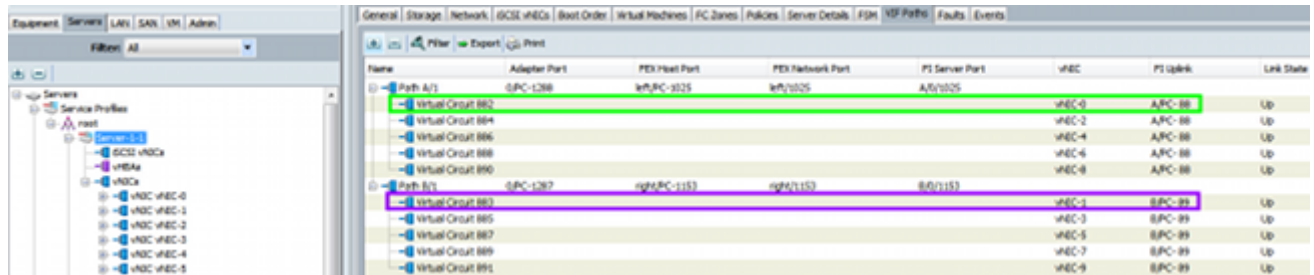
9. Confirm that the MAC address of the VMK interface

```
Mike-Cliff-Pod-16-B(nxos)# show mac address-table | in 8eb9
* 18      0050.5667.8eb9      dynamic  0          F    F    Veth883
Mike-Cliff-Pod-16-B(nxos)#
Mike-Cliff-Pod-16-B(nxos)# show int vethernet 883
Vethernet883 is up
  Bound Interface is port-channel1287
  Hardware: Virtual, address: 547f.eea3.c7e0 (bia 547f.eea3.c7e0)
  Description: server 1/1, VNIC vNIC-1
  Encapsulation ARPA
  Port mode is trunk
  EtherType is 0x8100
  Rx
    30553743 unicast packets  94871 multicast packets  1633080 broadcast packets
    32281694 input packets  32522468006 bytes
    0 input packet drops
  Tx
    16919347 unicast packets  588794 multicast packets  8994408 broadcast packets
    26502549 output packets  8364051391 bytes
    0 flood packets
    0 output packet drops
```

10. Check the pinning of these Veths to their uplinks with the **show circuit detail** command:

```
Mike-Cliff-Pod-16-B(nxos)# show circuit detail
Circuit Profile: Server-1
Serial: 1/1
  Veth-0
    Veth-0
    Link State: Up
    Oper State: Active
    Data Name:
    Admin Prio: 0
    Oper Prio: 0
    Mode: Trunk
    Transport: Ether
  Veth-1
    Veth-1
    Link State: Up
    Oper State: Active
    Data Name:
    Admin Prio: 0
    Oper Prio: 0
    Mode: Trunk
    Transport: Ether
```

**Note:** Other commands that output similar information are **show pinning server-interfaces**, **show pinning border-interfaces**, and **show pinning interface vethernet x**. You can also check the pinning in the UCSM:



11. Gather additional details about the port-channels. In this configuration, there are three port-channels in use for each FI. For example, FI-B has three associated port-channels:

- Port-channel 89 is the Link Aggregation Congrol Protocol (LACP) port-channel between FI-B and the upstream Nexus 5020.
- Port-channel 1153 is automatically created and is between FI-B and input/output module (IOM)-B.
- Port-channel 1287 is automatically created and is between IOM-B and Cisco VIC 1240 (blade).

1. Enter the **show port-channel summary** command in order to see the port-channel configuration of FI-B:

```
Mike-Cliff-Pod-16-B(nxos)# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
       I - Individual    H - Hot-standby (LACP only)
       s - Suspended     r - Module-removed
       S - Switched      R - Routed
       U - Up (port-channel)

-----
Group Port-      Type      Protocol  Member Ports
Channel
-----
89      Po89(SU)   Eth       LACP      Eth1/5(P)  Eth1/6(P)
1153    Po1153(SU) Eth       NONE      Eth1/1(P)
1287    Po1287(SU) Eth       NONE      Eth1/1/1(P) Eth1/1/3(P)
Mike-Cliff-Pod-16-B(nxos)#
```

2. Enter the **show cdp neighbors** command in order to discover and view additional information on FI-B:

```
Mike-Cliff-Pod-16-B(nxos)# 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
SJ-SV-C4K-1        mgmt0         179    R S I        WS-C4506      Gig5/40
N5K-Rack16-2 (FLC12110027) Eth1/5        163    S I s        N5K-C5020P-BA Eth1/22
N5K-Rack16-1 (SSI1351055H) Eth1/6        157    S I s        N5K-C5020P-BF Eth1/29
mc-vsm(1981308841355189719) Eth1/1/3      160    S I s        Nexus1000V    Eth5/2
```



- Enter the **show port-channel summary** command in order to see the port-channel configuration of FI-A:

```
Mike-Cliff-Pod-16-A(nxos)# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
       I - Individual    H - Hot-standby (LACP only)
       s - Suspended     r - Module-removed
       S - Switched      R - Routed
       U - Up (port-channel)
-----
Group Port-      Type      Protocol  Member Ports
Channel
-----
38    Po88 (SU)   Eth       LACP      Eth1/5 (P) Eth1/6 (P)
1025  Po1025 (SU) Eth       NONE      Eth1/1 (P)
1288  Po1288 (SU) Eth       NONE      Eth1/1/1 (P) Eth1/1/3 (P)
Mike-Cliff-Pod-16-A(nxos)#
```

- Enter the **show cdp neighbors** command in order to discover and view additional information on FI-A:

```
Mike-Cliff-Pod-16-A(nxos)# 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
SJ-SV-C4K-1        mgmt0          142    R S I      WS-C4506      Gig5/39
N5K-Rack16-2 (FLC12110027) Eth1/5         147    S I s      N5K-C5020P-BA Eth1/10
N5K-Rack16-1 (SSI1351055H) Eth1/6         121    S I s      N5K-C5020P-BF Eth1/11
mc-vsm(1981308841355189719) Eth1/1/1       167    S I s      Nexus1000V     Eth5/1
```

- Determine the specific pinning of the member interface from the port-channel.

Enter a **show port-channel** command in order to see that the FI-B - VMK interface

```
Mike-Cliff-Pod-16-B(nxos)# show port-channel load-balance forwarding-path interface port-channel 1287 vlan 18 src-mac 0050.5667.8eb9 dst-ip 172.16.18.1
Missing params will be substituted by 0's.
Load-balance Algorithm on FEK: source-dest-ip
crc8_hash: 209 Outgoing port id: Ethernet1/3
Param(s) used to calculate load-balance:
dst-ip: 172.16.18.1
src-ip: 0.0.0.0
dst-mac: 0000.0000.0000
src-mac: 0050.5667.8eb9
Mike-Cliff-Pod-16-B(nxos)#
Mike-Cliff-Pod-16-B(nxos)#
Mike-Cliff-Pod-16-B(nxos)# show port-channel load-balance forwarding-path interface port-channel 89 vlan 18 src-mac 0050.5667.8eb9 dst-ip 172.16.18.1
Missing params will be substituted by 0's.
Load-balance Algorithm on switch: source-dest-ip
crc8_hash: 5 Outgoing port id: Ethernet1/6
Param(s) used to calculate load-balance:
dst-ip: 172.16.18.1
src-ip: 0.0.0.0
dst-mac: 0000.0000.0000
src-mac: 0050.5667.8eb9
Mike-Cliff-Pod-16-B(nxos)#
```

Enter a **show port-channel** command in order to see that the FI-A - VM (ciscolive-vm) MAC address is pinned to Ethernet1/5 of port-channel 88:

```
Mike-Cliff-Pod-16-A(nxos)# show port-channel load-balance forwarding-path interface port-channel 38 vlan 18 src-mac 0000.5667.8eb9 dst-ip 172.16.18.1
Missing params will be substituted by 0's.
Load-balance Algorithm on FEK: source-dest-ip
crc8_hash: 209 Outgoing port id: Ethernet1/5
Param(s) used to calculate load-balance:
dst-ip: 172.16.18.1
src-ip: 0.0.0.0
dst-mac: 0000.0000.0000
src-mac: 0000.5667.8eb9
Mike-Cliff-Pod-16-A(nxos)#
Mike-Cliff-Pod-16-A(nxos)#
Mike-Cliff-Pod-16-A(nxos)# show port-channel load-balance forwarding-path interface port-channel 89 vlan 18 src-mac 0000.5667.8eb9 dst-ip 172.16.18.1
Missing params will be substituted by 0's.
Load-balance Algorithm on switch: source-dest-ip
crc8_hash: 5 Outgoing port id: Ethernet1/6
Param(s) used to calculate load-balance:
dst-ip: 172.16.18.1
src-ip: 0.0.0.0
dst-mac: 0000.0000.0000
src-mac: 0000.5667.8eb9
Mike-Cliff-Pod-16-A(nxos)#
```

13. Check that the MAC addresses are learned on the upstream Nexus 5020.

Enter a **show mac address-table** command in order to see that the VMK interface

```
N5K-Rack16-1#  
N5K-Rack16-1# show mac address-table | in 8eb9  
* 18      0050.5667.8eb9    dynamic    10          F      F      Po89  
N5K-Rack16-1#
```

Enter a **show mac address-table** command in order to see that the VM (ciscolive-vm) MAC address is learned on the Nexus 5020-2:

```
N5K-Rack16-2#  
N5K-Rack16-2# show mac address-table | in 63cc  
* 18      0050.568f.63cc    dynamic    0           F      F      Po88  
N5K-Rack16-2#
```

When you troubleshoot network issues, this example helps you quickly isolate and identify how and where a MAC address is learned and what the expected path is for network traffic.

## Verify

Verification procedures are included in the configuration example.

## Troubleshoot

This configuration example is intended to help with network troubleshooting.