



VI3 Networking: Advanced Troubleshooting

ESX Networking Architecture

Physical NICs

Virtual Switch

Virtual NICs


> VM's Virtual NIC

- Flexible
- Enhanced
- Virtual E1000

> Vswif for the Service Console

> VMkernel uses vmknics

VMkernel TCP/IP Stack



Agenda

Basic Troubleshooting

- > How to isolate problems?
- > What tools are available for troubleshooting?

Troubleshooting Scenarios

- > Step-by-step guide on how to troubleshoot some specific networking problems



Basic Troubleshooting Techniques


Isolate the problem

Troubleshoot one component at a time

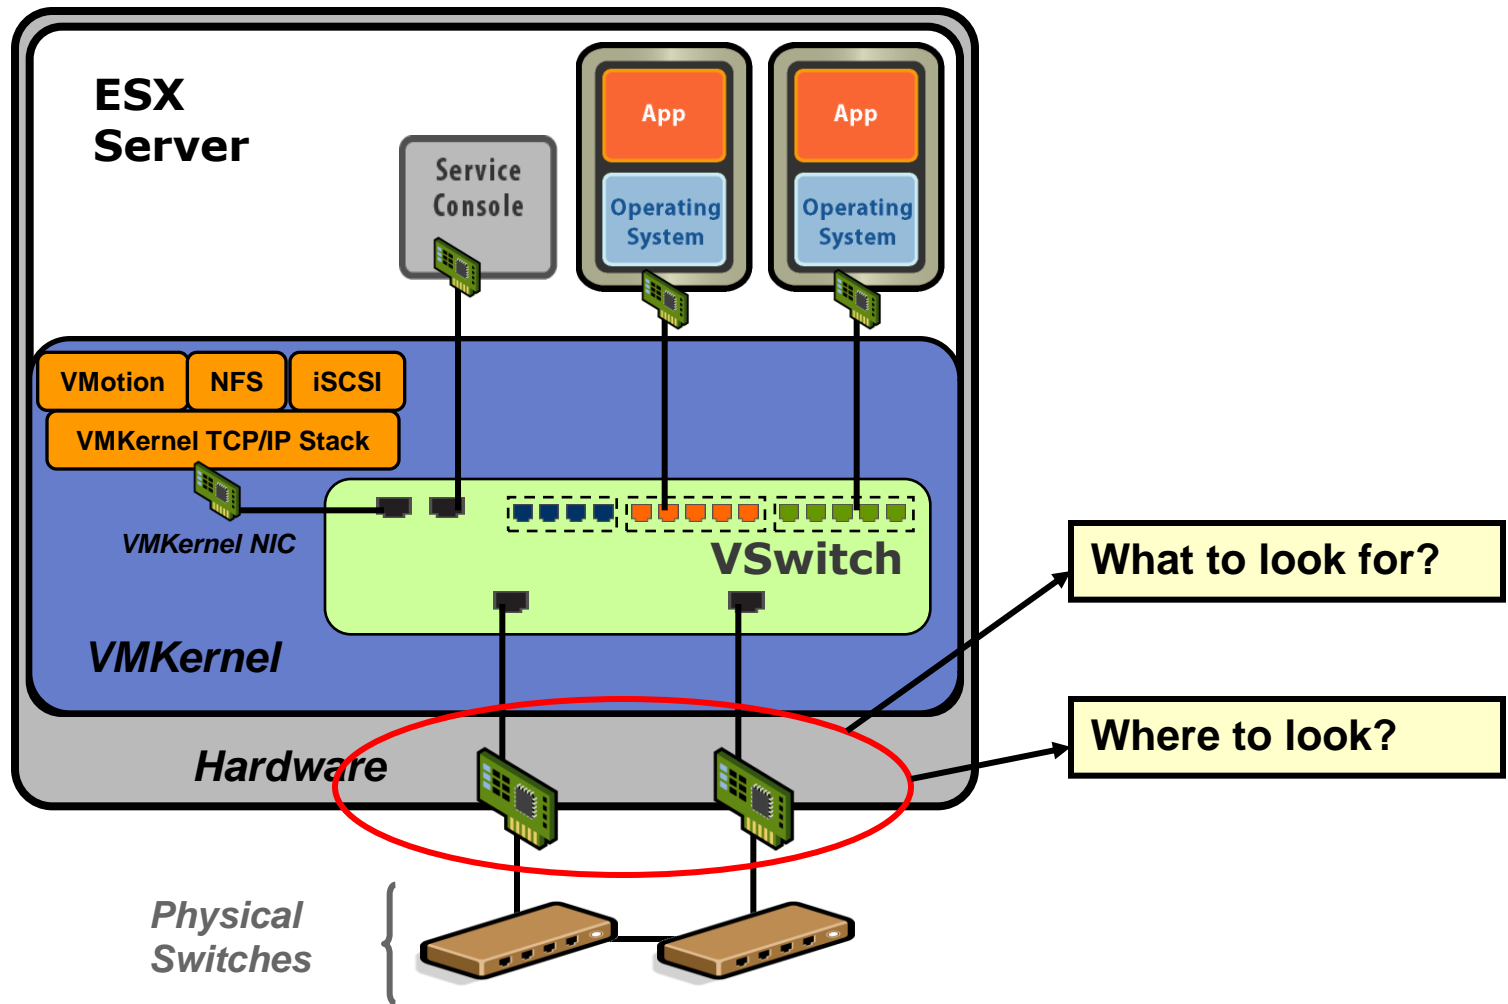
- > Physical NICs
- > Virtual Switch
- > Virtual NICs
- > Physical Network

Tools for troubleshooting

- > VI
- > Command Line Utilities
 - Ping and traceroute
 - Traffic sniffers and Protocol Analyzers
 - Wireshark
- > Logs



Isolating Network Problems: Physical NICs



Physical NICs: What to look for?

Does the device show up?

- > Is the driver loaded?

Physical properties of the link

- > Link State
- > Link Speed
- > Duplex Setting
- > MTU settings

Is the NIC connected to where you want it to be connected?

Is the NIC working?

- > Is the NIC transmitting and receiving packets?
- > Is the NIC dropping any packets?

esxcfg-nics
ifconfig inside
Service Console
VI Client

Network Hints
Cisco Discovery
Protocol (CDP)

ifconfig inside
Service Console
esxtop/resxtop
esxcfg-info

Getting information about the physical NICs

VI Client provides basic information about the physical NICs

Getting Started Summary Virtual Machines Resource Allocation Performance Configuration Users & Groups Events Permissions

Hardware

- Health Status
- Processors
- Memory
- Storage
- Networking
- Storage Adapters
- Network Adapters

Network Adapters

Device	Speed	Configured	vSwitch	Observed IP ranges	Wake on LAN Supported
NC7781 Gigabit Server Adapter (PCI-X, 10,100,1000-T)					
vmnic0	100 Full	Negotiate	vSwitch0	10.17.40.1-10.17.43.254	Yes
vmnic1	100 Full	Negotiate	None	10.17.40.1-10.17.43.254	Yes
8254NXX Gigabit Ethernet Controller					
vmnic2	100 Full	Negotiate	None	10.17.40.1-10.17.43.254	Yes
EtherExpress PRO/100 S Server Adapter					
vmnic3	100 Full	Negotiate	None	10.17.40.1-10.17.43.254	Yes
vmnic4	100 Full	Negotiate	None	10.17.40.1-10.17.43.254	Yes

Type of NIC

Link Status

Connections

Network Hint

Getting information about the physical NICs

esxcfg-nics allows you to set or get physical NIC settings via the command line

```
[root@osdc-sust005 root]# esxcfg-nics
esxcfg-nics <options> [nic]
-s|--speed <speed>      Set the speed of this NIC to one of 10/100/1000/10000.
                          Requires a NIC parameter.
-d|--duplex <duplex>    Set the duplex of this NIC to one of 'full' or 'half'.
                          Requires a NIC parameter.
-a|--auto                Set speed and duplexity automatically. Requires a NIC parameter.
-l|--list                Print the list of NICs and their settings.
-r|--restore             Restore the nics configured speed/duplex settings (INTERNAL ONLY)
-h|--help               Display this message.
```

Getting information about the physical NICs

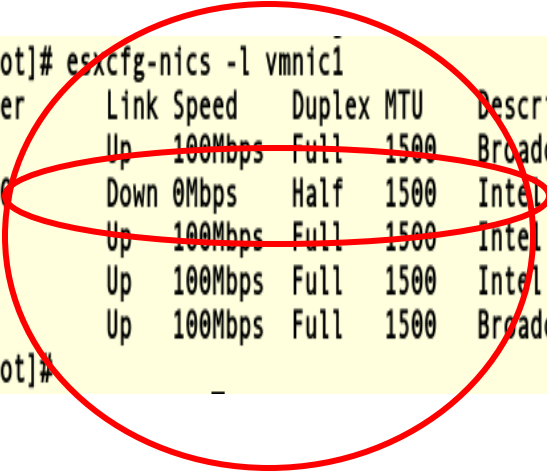
The '- l' option lists the nics in the system and their settings

- > Link State
- > Speed
- > Duplex
- > MTU

```
[root@osdc-sust005 root]# esxcfg-nics -l vmnic1
```

Name	PCI	Driver	Link	Speed	Duplex	MTU	Description
vmnic1	01:05.00	tg3	Up	100Mbps	Full	1500	Broadcom Corporation NC7781 Gigabit Server Adapter (PCI-X, 10,100,1000-T)
vmnic2	07:02.00	e1000	Down	0Mbps	Half	1500	Intel Corporation 8254NXX Gigabit Ethernet Controller
vmnic4	0b:02.00	e1000	Up	100Mbps	Full	1500	Intel Corporation EtherExpress PRO/100 S Server Adapter
vmnic3	0b:01.00	e1000	Up	100Mbps	Full	1500	Intel Corporation EtherExpress PRO/100 S Server Adapter
vmnic0	01:04.00	tg3	Up	100Mbps	Full	1500	Broadcom Corporation NC7781 Gigabit Server Adapter (PCI-X, 10,100,1000-T)

```
[root@osdc-sust005 root]#
```



Where is the physical NIC connected to?

Just follow the cable

OR


Use CDP and Network Hints



Cisco Discovery Protocol

Periodic exchange of information

- > Physical switch port a vmnic is connected to
- > vSwitch a physical switch port is connected to
- > Duplex and speed settings



Cisco Discovery Protocol	
Properties	
Version	0
Timeout	0
Time to live	141
Samples	11290
Device Id	cs-eth-netcore1.csl.vmwar
Address	172.17.1.3
Port Id	GigabitEthernet0/23
Software Version	Cisco IOS Software, C3560
Hardware Platform	cisco WS-C3560G-24TS
IP Prefix	0.0.0.0
IP Prefix Length	0
VLAN	1
Full Duplex	true
MTU	0
System Name	
System OId	
Management Address	172.17.1.3
Location	
CDP Device Capability	
Router	false
Transparent Bridge	false
Source Route Bridge	false
Network Switch	true
Host	false
IGMP Enabled	true
Repeater	false

Cisco Discovery Protocol

CDP is enabled by default in listening mode

On ESX Server 3.5, it is possible to configure CDP also in advertising mode

> Enabled/disabled only via command line with

esxcfg-vswitch -B <state> <vSwitch>

■ States

- Listen
- Advertise
- Both
- Down

> Verify the setting with

esxcfg-vswitch -b <vSwitch>

Getting information about the physical NICs

esxtop provides system-wide real-time traffic information

➤ For ESXi use resxtop utility provided in the RCLI

Type 'n' to switch to the network utilization screen

Output of esxtop

11:45:05pm up 5 min, 63 worlds; CPU load average: 0.00, 0.04, 0.00

PORT-ID	USED-BY	TEAM-PNIC	DNAME	PKTTX/s	MbTX/s	PKTRX/s	MbRX/s	%DRPTX	%DRPRX
16777217	vmnic0	-	vSwitch0	1.98	0.00	15.05	0.01	0.00	0.00
16777218	0:NCP	-	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777219	0:CDP	-	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777220	0:vswif0	vmnic0	vSwitch0	1.98	0.00	14.26	0.01	0.00	0.00
33554433	vmnic3	-	vSwitch1	0.00	0.00	13.27	0.01	0.00	0.00
33554434	0:NCP	-	vSwitch1	0.00	0.00	0.00	0.00	0.00	0.00
33554435	vmnic4	-	vSwitch1	1597.13	1.19	1610.20	1.20	0.00	0.00
33554436	0:CDP	-	vSwitch1	0.00	0.00	0.00	0.00	0.00	0.00
33554437	0:vswif10	vmnic3	vSwitch1	0.00	0.00	12.28	0.01	0.00	0.00
33554438	0:vmk-tcpip-10.17.41	vmnic4	vSwitch1	1597.13	1.19	1609.41	1.20	0.00	0.00

Getting information about the physical NICs

On ESX Server 3.5, running `ifconfig` inside Service Console provides information valuable for troubleshooting

Output of ifconfig

```
[root@osdc-sust005 root]# ifconfig
lo      Link encap:Local Loopback
        inet addr:127.0.0.1  Mask:255.0.0.0
        UP LOOPBACK RUNNING  MTU:16436  Metric:1
        RX packets:764267 errors:0 dropped:0 overruns:0 frame:0
        TX packets:764267 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:559614933 (533.6 Mb)  TX bytes:559614933 (533.6 Mb)

vmnic0  Link encap:Ethernet  HWaddr 00:11:85:13:79:9A
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:9235676 errors:0 dropped:0 overruns:0 frame:0
        TX packets:410441 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:666112677 (635.2 Mb)  TX bytes:202118997 (192.7 Mb)
        Interrupt:153

vmnic1  Link encap:Ethernet  HWaddr 00:11:85:13:79:9B
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:8989025 errors:0 dropped:0 overruns:0 frame:0
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:604831816 (576.8 Mb)  TX bytes:0 (0.0 b)
        Interrupt:161
```


Getting information about the physical NICs

esxcfg-info provides detailed information about the system

- > Use '*esxcfg-info -n*' for network information
- > Redirect the output of esxcfg-info to a file
- > Look for the 'Physical Nic' section

```
[root@osdc-sust005 root]# esxcfg-info > info
[root@osdc-sust005 root]# vi info
[root@osdc-sust005 root]# █
```


Getting information about the physical NICs

```
\==+Physical Nic :
|---Name.....vmnic0
|---PCI Bus.....1
|---PCI Slot.....4
|---PCI function.....0
|---MAC Address.....00:11:85:13:79:9a
|---Virtual MAC Address.....00:50:56:53:79:9a
|---Driver.....tg3
|---Network Hint.....0 10.17.40.00/255.255.254.00
|---MTU.....1500
|---Configured Speed.....0
|---Actual Speed.....100
|---Configured Duplex.....0
|---Actual Duplex.....1
|---Link Up.....true
|---Wake on LAN supported .....true
|---Wake on LAN enabled .....true
|---CDP Network.....false
\==+Hardware Capabilities :
|---Vlan Tag.....true
|---Vlan Untag.....true
|---Checksum Offload.....true
|---Tcp Segmentation Offload.....false
|---Scatter Gather Tx.....true
\==+Supported Speeds and Duplex :
|---Supported Mode.....auto/auto
|---Supported Mode.....10/half
|---Supported Mode.....10/full
|---Supported Mode.....100/half
|---Supported Mode.....100/full
|---Supported Mode.....1000/half
|---Supported Mode.....1000/full
\==+PCI Device :
|---Bus.....0x01
|---Slot.....0x04
|---Function.....0x00
|---Vendor Id.....0x14e4
|---Device Id.....0x16a7
|---Sub-Vendor Id.....0x0e11
|---Sub-Device Id.....0x00cb
|---Vendor Name.....Broadcom Corporation
|---Device Name.....NC7781 Gigabit Server Adapter (PCI-X, 10,100,1000-T)
|---Device Class.....512
|---PIC Line.....3
|---Service Console IRQ.....38
|---Vector.....153
```

Output of esxcfg-info

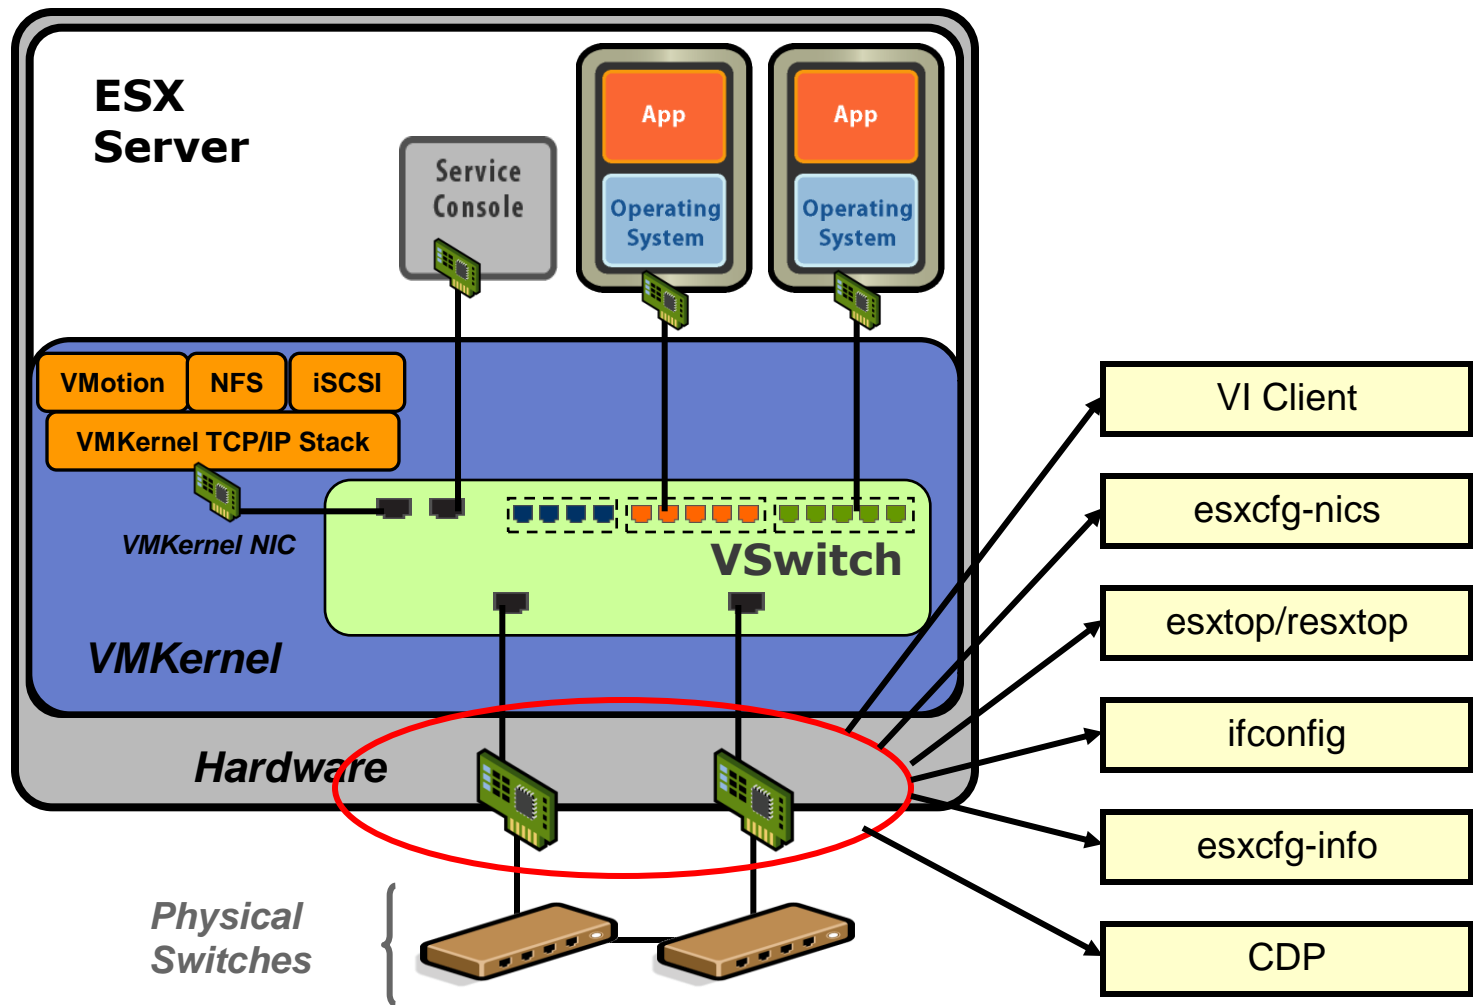
Getting information about the physical NICs

Search for vmnicX in the output of esxcfg-info

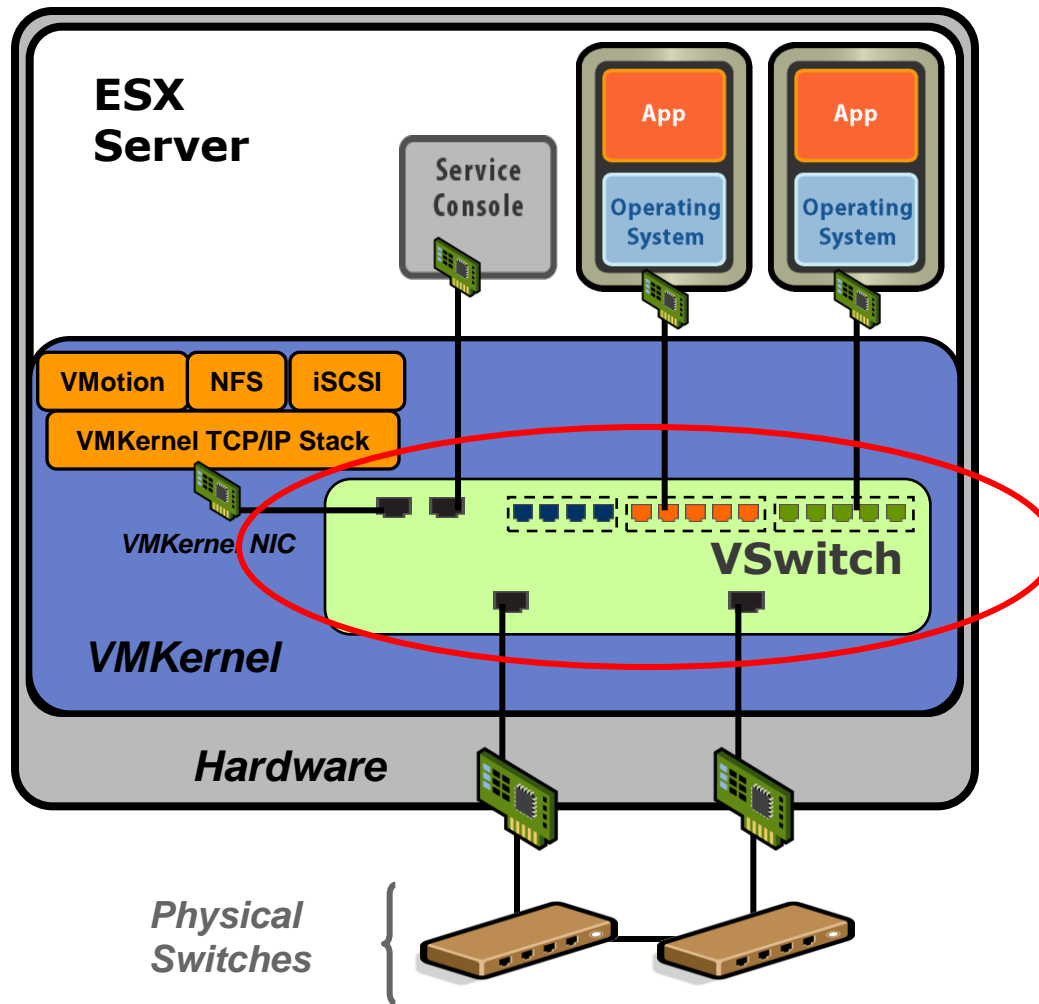
```
\==+Virtual Switch :
|---Name.....vSwitch0
|---Uplinks.....vmnic0
|---Total Ports.....64
|---Configured Ports.....64
|---Used Ports.....4
|---Beacon Enabled.....false
|---Beacon Interval.....1
|---Beacon Threshold.....1
|---Beacon Required By.....
|---CDP Status.....listen
|---MTU.....1500
|---Switch Class.....etherswitch
\==+Unclaimed Ports :
  \==+Port :
    |---Port Id.....16777217
    |---World Leader.....0
    |---Client Name.....vmnic0
    |---MAC Addr.....00:00:00:00:00:00
    |---Blocked.....false
    |---Type.....Pnic
  \==+Stats :
    |---Packets Tx Ok.....410963
    |---Bytes Tx Ok.....202209073
    |---Dropped Tx.....0
    |---Packets TSO Tx Ok.....0
    |---Bytes TSO Tx Ok.....0
    |---Dropped TSO Tx.....0
    |---Packets SW TSO Tx.....0
    |---Dropped SW TSO Tx.....0
    |---Packets Zero Copy Tx Ok.....0
    |---Packets Rx Ok.....9242292
    |---Bytes Rx Ok.....666574263
    |---Dropped Rx.....0
```

Output of esxcfg-info

Summary: Getting information about the physical NICs



Isolating Network Problems: Virtual Switch



Virtual Switch: What to look for?

vSwitch and Portgroup Configuration

- > Uplinks
- > VLAN Setting
- > Layer 2 Security Policies
- > NIC Teaming Configuration

esxcfg-vswitch
esxcfg-info
VI Client

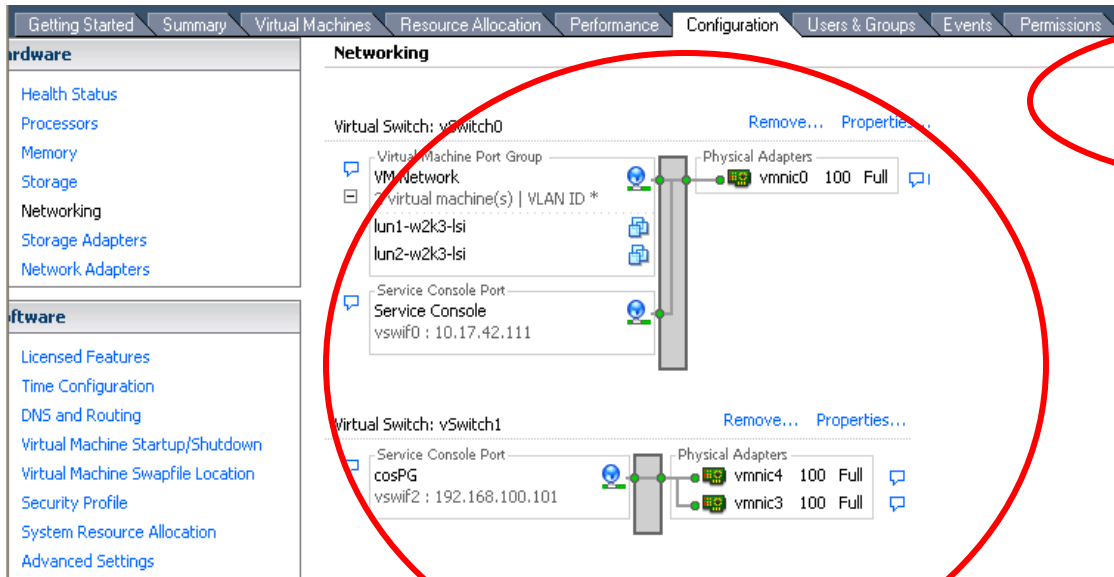
Is the traffic flowing through the vSwitch?

- > Is the vSwitch dropping any packets?

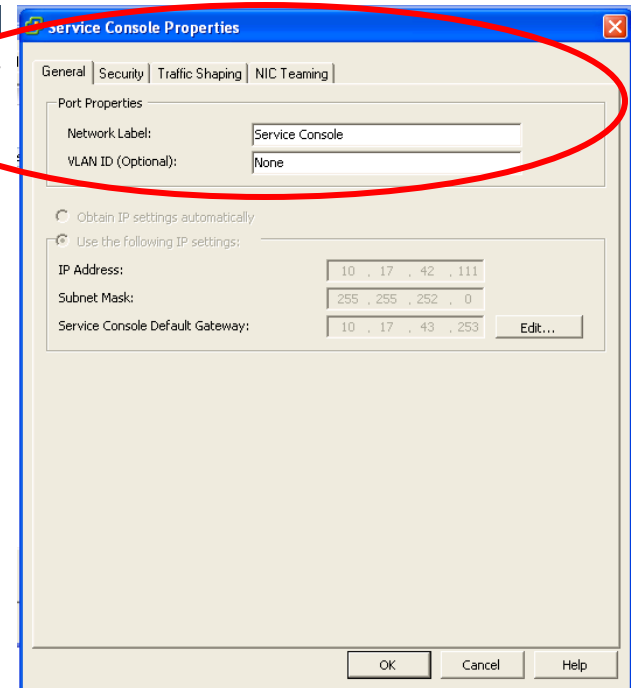
esxtop/resxtop
esxcfg-info

Getting information about the vSwitch: VI

VI: Virtual Switch Configuration



VI: Portgroup Properties



Getting information about the vSwitch: esxcfg-vswitch

esxcfg-vswitch

- An interface for adding, removing, and modifying virtual switches and their settings

Output of esxcfg-vswitch -l

```
[root@osdc-sust005 root]# esxcfg-vswitch -l
```

Switch Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
vSwitch0	64	4	64	1500	vmnic0
PortGroup Name	VLAN ID	Used Ports	Uplinks		
VM Network	0	0	vmnic0		
Service Console	0	1	vmnic0		

Switch Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
vSwitch1	64	4	64	1500	vmnic4,vmnic3
PortGroup Name	VLAN ID	Used Ports	Uplinks		
cosPG	0	0	vmnic3,vmnic4		

Getting information about the vSwitch: esxtop

esxtop provides system-wide real-time traffic information

➤ For ESXi use resxtop utility provided in the RCLI

Type 'n' to switch to the network utilization screen


Output of esxtop

11:45:05pm up 5 min, 63 worlds; CPU load average: 0.00, 0.04, 0.00

PORT-ID	USED-BY	TEAM-PNIC	DNAME	PKTTX/s	MbTX/s	PKTRX/s	MbRX/s	%DRPTX	%DRPRX
16777217	vmnic0	-	vSwitch0	1.98	0.00	15.05	0.01	0.00	0.00
16777218	0:NCP	-	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777219	0:CDP	-	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777220	0:vswif0	vmnic0	vSwitch0	1.98	0.00	14.26	0.01	0.00	0.00
33554433	vmnic3	-	vSwitch1	0.00	0.00	13.27	0.01	0.00	0.00
33554434	0:NCP	-	vSwitch1	0.00	0.00	0.00	0.00	0.00	0.00
33554435	vmnic4	-	vSwitch1	1597.13	1.19	1610.20	1.20	0.00	0.00
33554436	0:CDP	-	vSwitch1	0.00	0.00	0.00	0.00	0.00	0.00
33554437	0:vswif10	vmnic3	vSwitch1	0.00	0.00	12.28	0.01	0.00	0.00
33554438	0:vmk-tcpip-10.17.41	vmnic4	vSwitch1	1597.13	1.19	1609.41	1.20	0.00	0.00

Real Time Traffic
Information

Getting information about the vSwitch: esxtop



vmknic traffic is going through vmnic4

Output of esxtop

11:45:05pm up 5 min, 63 worlds; CPU load average: 0.00 0.04, 0.00

PORT-ID	USED-BY	TEAM-PNIC	ONAME	PKTTX/s	MbTX/s	PKTRX/s	MbRX/s	%DRPTX	%DRPRX
16777217	vmnic0	-	vSwitch0	1.98	0.00	15.05	0.01	0.00	0.00
16777218	0:NCP	-	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777219	0:CDP	-	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777220	0:vswif0	vmnic0	vSwitch0	1.98	0.00	14.26	0.01	0.00	0.00
33554433	vmnic3	-	vSwitch1	0.00	0.00	13.27	0.01	0.00	0.00
33554434	0:NCP	-	vSwitch1	0.00	0.00	0.00	0.00	0.00	0.00
33554435	vmnic4	-	vSwitch1	1597.13	1.19	1610.20	1.20	0.00	0.00
33554436	0:CDP	-	vSwitch1	0.00	0.00	0.00	0.00	0.00	0.00
33554437	0:vswif10	vmnic3	vSwitch1	0.00	0.00	12.28	0.01	0.00	0.00
33554438	0:vmk-tcpip-10.17.41	vmnic4	vSwitch1	1597.13	1.19	1609.41	1.20	0.00	0.00

Getting information about the vSwitch: esxcfg-info

esxcfg-info provides information in greater detail

> Configuration information

Output of esxcfg-info

```
\==+Virtual Switch :
|---Name.....vSwitch1
|---Uplinks.....vmnic4,vmnic3
|---Total Ports.....64
|---Configured Ports.....64
|---Used Ports.....6
|---Beacon Enabled.....false
|---Beacon Interval.....1
|---Beacon Threshold.....1
|---Beacon Required By.....
|---CDP Status.....listen
|---MTU.....1500
|---Switch Class.....etherswitch
```

Getting information about the vSwitch: esxcfg-info

esxcfg-info provides information in greater detail

➤ Cumulative traffic information for each port on the vSwitch

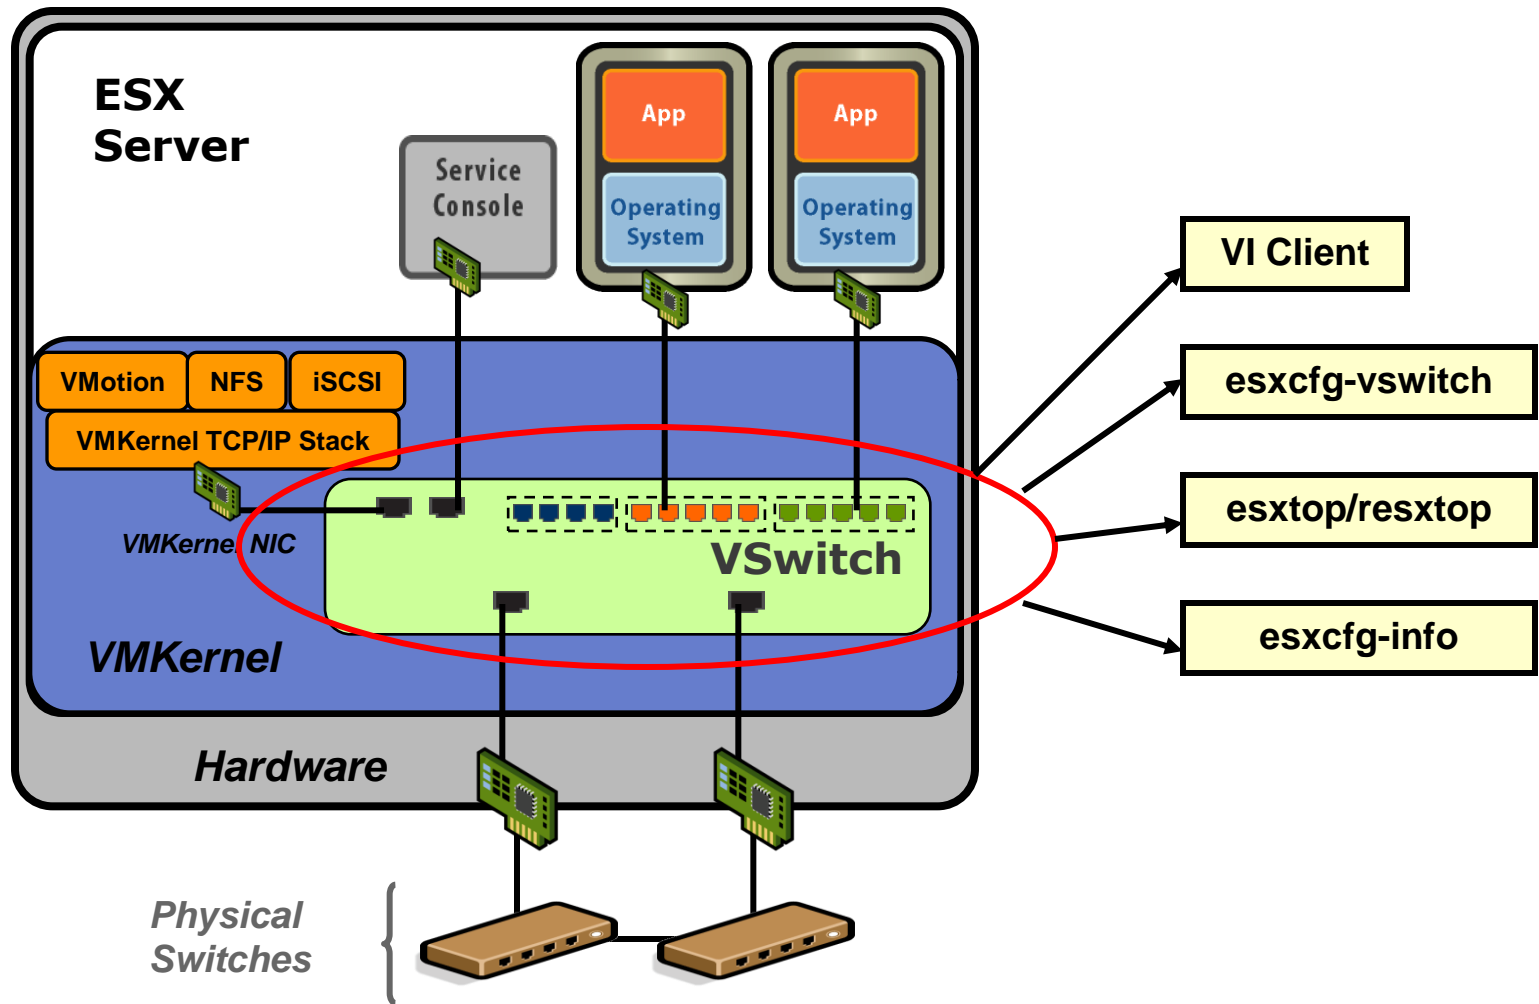
```
\==+Port :
|----Port Id.....33554438
|----World Leader.....0
|----Client Name.....vmk-tcpip-10.17.41.72
|----MAC Addr.....00:50:56:7e:31:61
|----Blocked.....false
|----Type.....Tcp/Ip
|----Portgroup Name.....vmkPG
\==+Stats :
|----Packets Tx Ok.....84489
|----Bytes Tx Ok.....8281172
|----Dropped Tx.....0
|----Packets TSO Tx Ok.....0
|----Bytes TSO Tx Ok.....0
|----Dropped TSO Tx.....0
|----Packets SW TSO Tx.....0
|----Dropped SW TSO Tx.....0
|----Packets Zero Copy Tx Ok.....0
|----Packets Rx Ok.....113655
|----Bytes Rx Ok.....10153132
|----Dropped Rx.....0
|----Dropped TSO Rx.....0
|----Packets SW TSO Rx.....0
|----Dropped SW TSO Rx.....0
|----Actions.....113495
|----Uplink Rx Packets.....142780
|----Pks Billed.....24406
|----Dropped Tx Due to Page Absent....0
|----Dropped Rx Due to Page Absent....0
```

**Information about VMkernel
Port**


```
\==+Port :
|----Port Id.....33554435
|----World Leader.....0
|----Client Name.....vmnic4
|----MAC Addr.....00:00:00:00:00:00
|----Blocked.....false
|----Type.....Pnic
\==+Stats :
|----Packets Tx Ok.....84488
|----Bytes Tx Ok.....8281112
|----Dropped Tx.....0
|----Packets TSO Tx Ok.....0
|----Bytes TSO Tx Ok.....0
|----Dropped TSO Tx.....0
|----Packets SW TSO Tx.....0
|----Dropped SW TSO Tx.....0
|----Packets Zero Copy Tx Ok.....0
|----Packets Rx Ok.....118277
|----Bytes Rx Ok.....10538216
|----Dropped Rx.....0
|----Dropped TSO Rx.....0
|----Packets SW TSO Rx.....0
|----Dropped SW TSO Rx.....0
|----Actions.....0
|----Uplink Rx Packets.....0
|----Pks Billed.....0
|----Dropped Tx Due to Page Absent....0
|----Dropped Rx Due to Page Absent....0
```

**Information about Uplink Port
(vmnic4)**

Summary: Getting information about the vSwitch



Isolating Network Problems: Virtual NICs



Virtual NICs: What to look for?

Does the device show up?

- > Is the driver loaded?

Physical properties of the link

- > Link State
- > MTU settings

Is the vNIC connected to the correct portgroup?

- > Portgroup using the correct uplink
- > Portgroup with the correct security properties

Is the NIC working?

- > Does the NIC have an IP address?
- > Is the NIC transmitting and receiving packets?
- > Is the NIC dropping any packets?

esxcfg-vswif

esxcfg-vmknics

Guest specific utilities

> Linux

- ifconfig
- lspci

> Windows

- Device Manager

VI Client

.vmx file

esxcfg-info

Guest specific utilities

> Linux

- ifconfig

> Windows

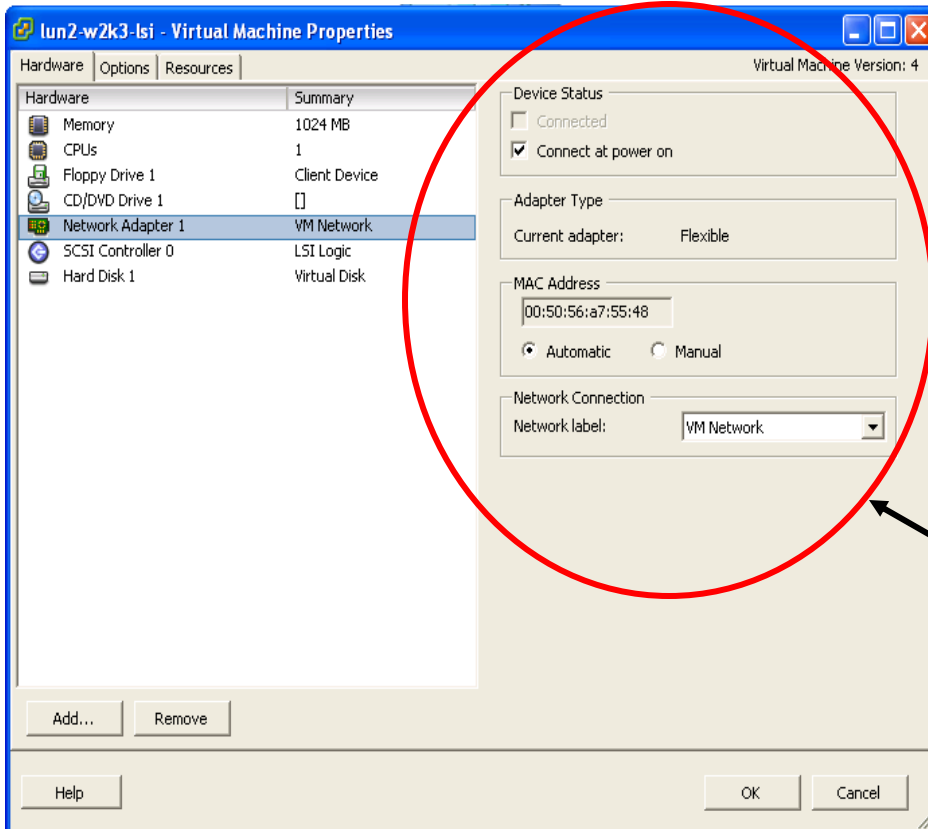
- Network Connections

esxtop/resxtop

esxcfg-info

Getting information about the vNIC

VI: Virtual Machine Properties



.vmx file

```
ethernet0.present = "true"  
ethernet0.allowGuestConnectionControl = "false"  
ethernet0.networkName = "VM Network"  
ethernet0.addressType = "vpx"  
ethernet0.generatedAddress = "00:50:56:a7:55:48"
```

VM's Connection Information

Getting information about the vNIC

esxcfg-vswif

- > An interface to configure Service Console NIC

esxcfg-vmknic

- > An interface to configure VMkernel NIC

Output of esxcfg-vswif -l

```
[root@osdc-sust005 root]# esxcfg-vswif -l
```

Name	Port Group	IP Address	Netmask	Broadcast	Enabled	DHCP
vswif0	Service Console	10.17.42.111	255.255.252.0	10.17.43.255	true	false
vswif10	cosPG	10.17.40.156	255.255.252.0	10.17.43.255	true	true

Output of esxcfg-vmknic -l

```
[root@osdc-sust005 root]# esxcfg-vmknic -l
```

Interface	Port Group	IP Address	Netmask	Broadcast	MAC Address	MTU	TSO MSS	Enabled
vmk0	vmkPG	10.17.41.72	255.255.252.0	10.17.43.255	00:50:56:7e:31:61	1500	40960	true

Getting information about the vNIC

Output of esxtop

11:45:05pm up 5 min, 63 worlds; CPU load average: 0.00, 0.04, 0.00

PORT-ID	USED-BY	TEAM-PNIC	DNAME	PKTTX/s	MbTX/s	PKTRX/s	MbRX/s	%DRPTX	%DRPRX
16777217	vmnic0	-	vSwitch0	1.98	0.00	15.05	0.01	0.00	0.00
16777218	0:NCP	-	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777219	0:CDP	-	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777220	0:vswif0	vmnic0	vSwitch0	1.98	0.00	14.26	0.01	0.00	0.00
33554433	vmnic3	-	vSwitch1	0.00	0.00	13.27	0.01	0.00	0.00
33554434	0:NCP	-	vSwitch1	0.00	0.00	0.00	0.00	0.00	0.00
33554435	vmnic4	-	vSwitch1	1597.13	1.19	1610.20	1.20	0.00	0.00
33554436	0:CDP	-	vSwitch1	0.00	0.00	0.00	0.00	0.00	0.00
33554437	0:vswif10	vmnic3	vSwitch1	0.00	0.00	12.28	0.01	0.00	0.00
33554438	0:vmk-tcpip-10.17.41	vmnic4	vSwitch1	1597.13	1.19	1609.41	1.20	0.00	0.00

Output of esxcfg-info

```
\==+Port :
|----Port Id.....33554438
|----World Leader.....0
|----Client Name.....vmk-tcpip-10.17.41.72
|----MAC Addr.....00:50:56:7e:31:61
|----Blocked.....false
|----Type.....Tcp/Ip
|----Portgroup Name.....vmkPG
\==+Stats :
|----Packets Tx Ok.....84489
|----Bytes Tx Ok.....8281172
|----Dropped Tx.....0
|----Packets TSO Tx Ok.....0
|----Bytes TSO Tx Ok.....0
|----Dropped TSO Tx.....0
|----Packets SW TSO Tx.....0
|----Dropped SW TSO Tx.....0
|----Packets Zero Copy Tx Ok.....0
|----Packets Rx Ok.....113655
|----Bytes Rx Ok.....10153132
```

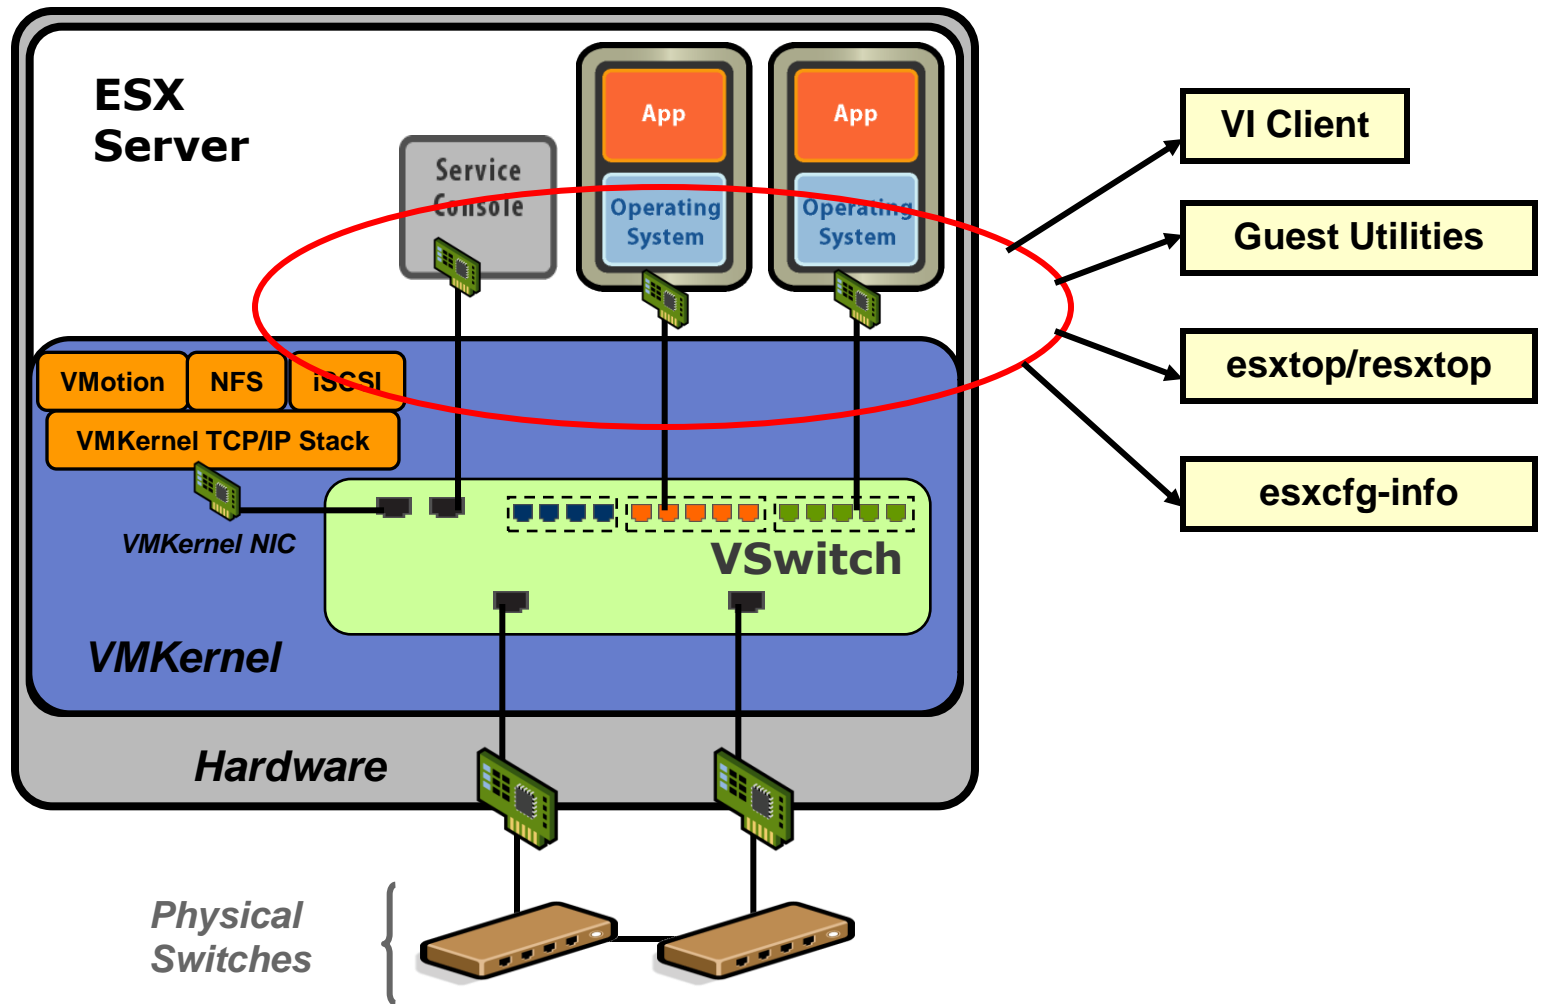
Real time traffic information

Cumulative Traffic Information

Look for Rx/Tx information for the vNIC you are interested in

Search for the port ID of the vNIC in the esxcfg-info output

Summary: Getting information about the vNIC



Sniffing For Trouble

Sniff for packets at different layers for isolation

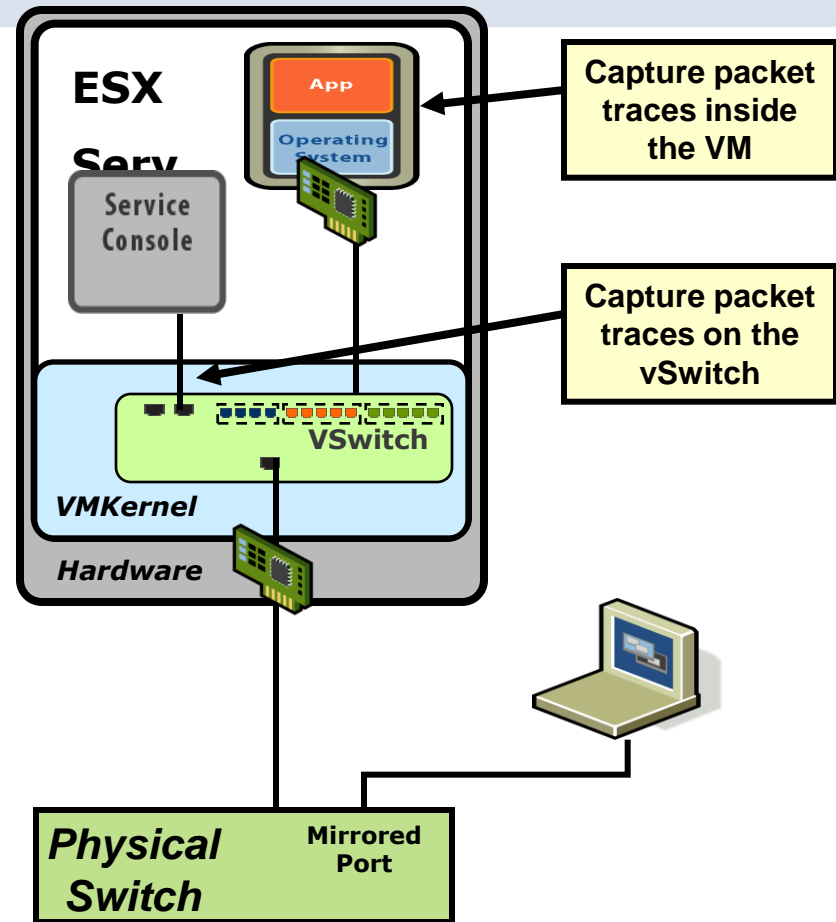
- > Physical Switch Port Level
- > vSwitch Level
- > VM Level

Look for

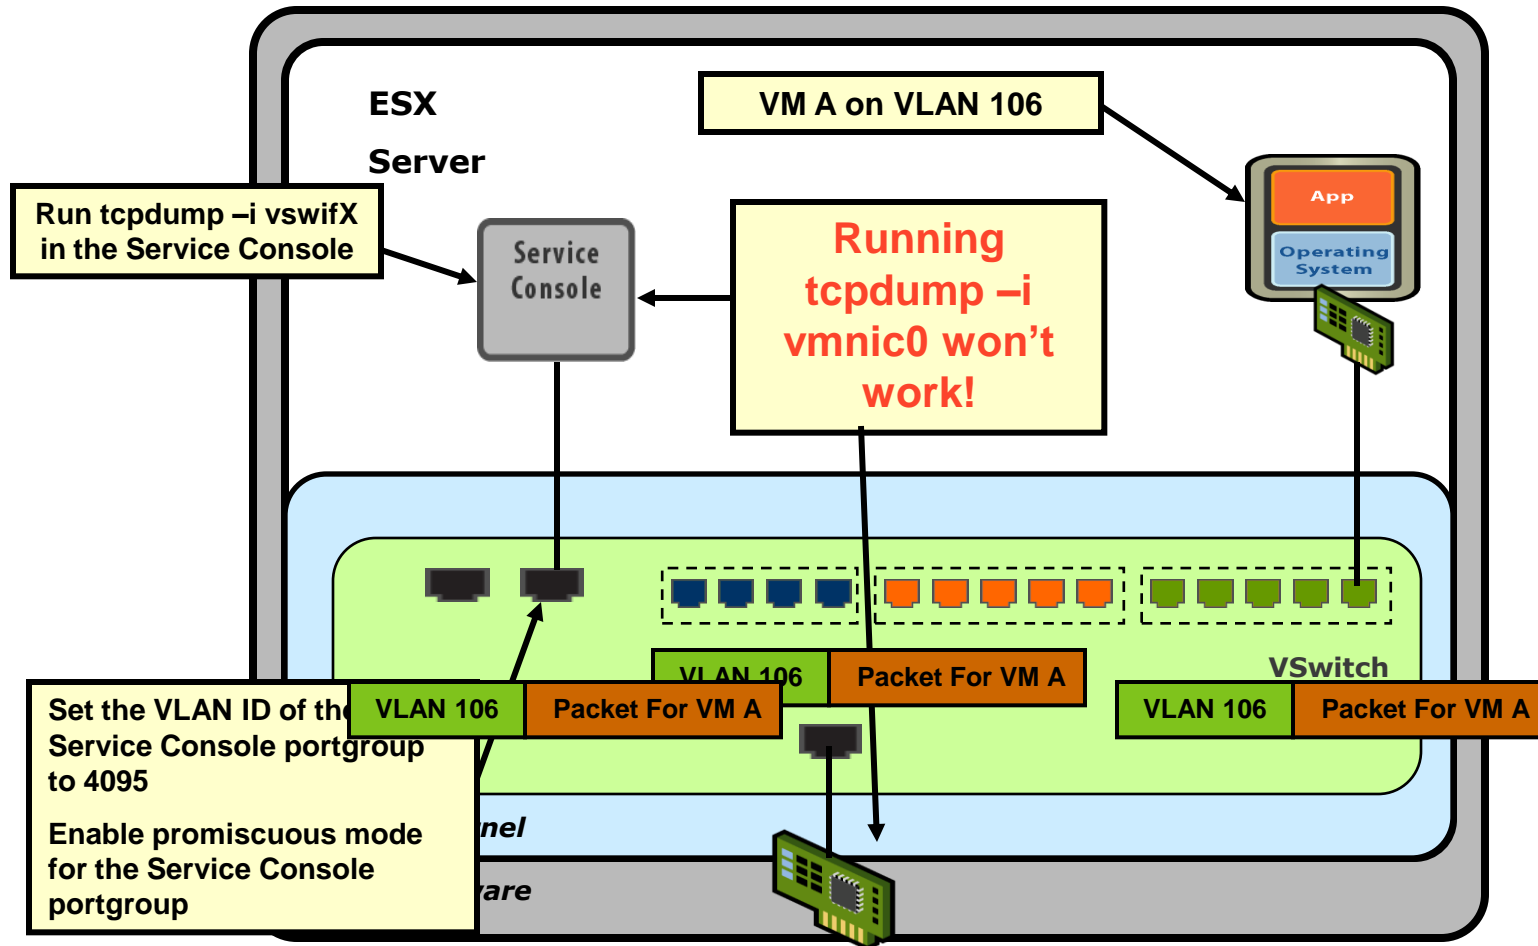
- > Lost Packets
- > Large number of packet retransmissions
- > Anomalies reported by protocol analyzers like Wireshark etc.

Look for patterns

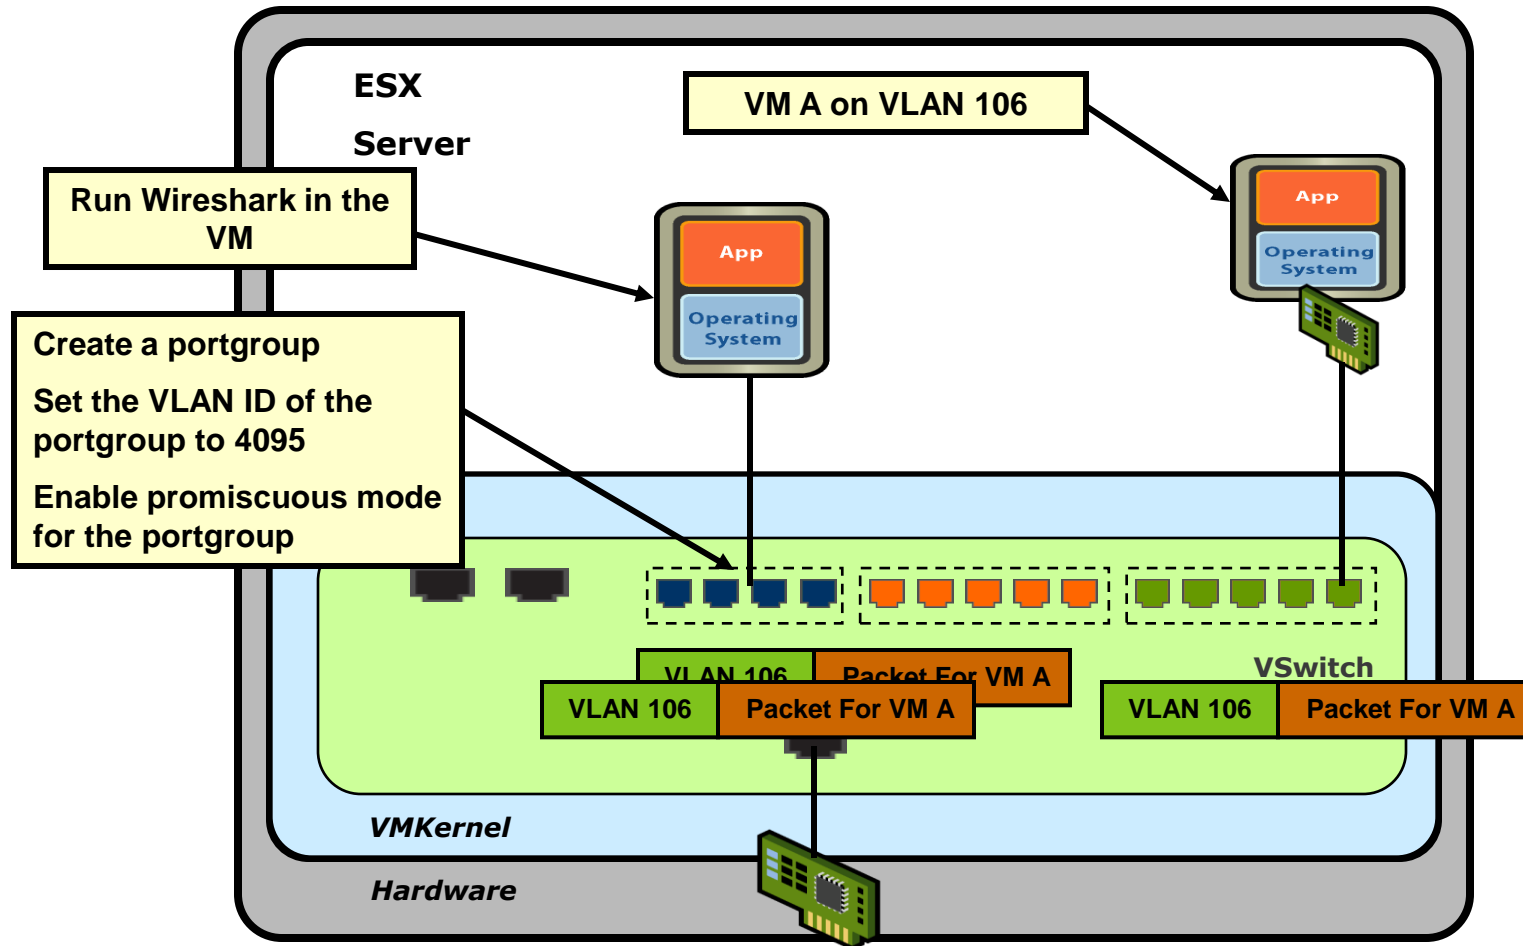
- > Are packets of a certain type causing problems?
- > Are packets of a certain size causing problems?



Collecting Network Traces on the vSwitch



Collecting Network Traces on the vSwitch



Logs on ESX

VMkernel logs

- > `/var/log/vmkernel` for ESX
- > `/var/log/messages` for ESXi

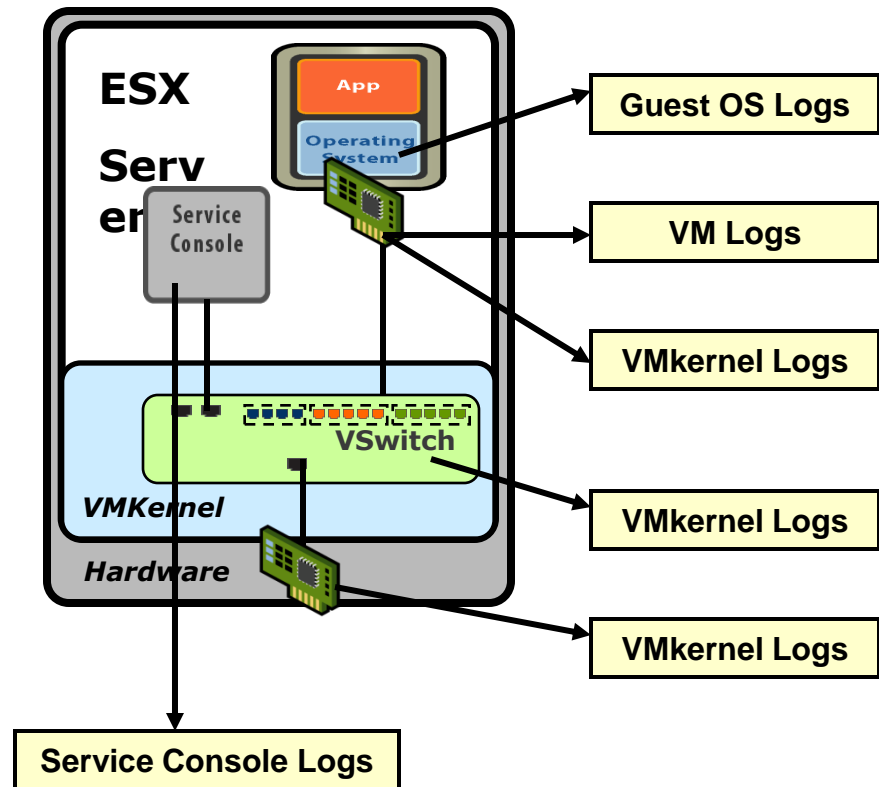
VM logs

- > `vmware.log` file in the VM directory

Service console logs

- > `/var/log/messages` for ESX

Also check the guest OS logs for any errors





Troubleshooting Scenarios

Signs of trouble

Basic connectivity problems

- > No network connectivity on some or all of the VMs on a vSwitch

Flaky network connection

- > Connection timeouts
- > Intermittent loss of connectivity

Problem

None of the VMs on my ESX box have network connectivity

Step 1: Check inside the VM

Has the guest OS detected the vNIC?

- > Use OS specific utilities to check
 - Windows: Device Manager
 - Linux: lspci, ifconfig
- > If a network interface for the vNIC does not show up
 - Check the VM configuration using VI or by looking into the .vmx file
 - Check the VM's log file (*vmware.log*) for any obvious problems
 - Check the guest OS log files for any obvious problems
 - Is appropriate driver for the vNIC installed and loaded?
 - Install VMware tools inside the guest

Step 1: Check inside the VM

Is the network interface in the guest up and does it have an IP address?

- > Use OS specific utilities to check
 - Windows: Network connections, ipconfig
 - Linux: ifconfig
- > Use static IP addresses during troubleshooting

Step 2: Check the vNIC connection

Check if the vNIC is connected to the correct portgroup

- > Use VI or look into the .vmx file
- > Make sure the 'Connected' box is checked

VI: Virtual Machine Properties

The screenshot shows the 'Virtual Machine Properties' dialog box for a VM named 'lun2-w2k3-lsi'. The 'Hardware' tab is selected, and 'Network Adapter 1' is highlighted in the left pane. The right pane shows the configuration for this adapter. Two red circles highlight specific settings, each with a callout box:

- The first circle is around the 'Device Status' section, where the 'Connect at power on' checkbox is checked. A callout box points to it with the text: "Is the vNIC connected".
- The second circle is around the 'Network Connection' section, where the 'Network label' dropdown is set to 'VM Network'. A callout box points to it with the text: "The vNIC connects to this portgroup".

Hardware	Summary
Memory	1024 MB
CPUs	1
Floppy Drive 1	Client Device
CD/DVD Drive 1	[]
Network Adapter 1	VM Network
SCSI Controller 0	LSI Logic
Hard Disk 1	Virtual Disk

Virtual Machine Version: 4

Device Status

- ☐ Connected
- ☒ Connect at power on

Adapter Type

Current adapter: Flexible

MAC Address

00:50:56:a7:55:48

☒ Automatic ☐ Manual


Network Connection

Network label: VM Network

Step 2: Check the vNIC connection

Check connectivity between VMs on the same portgroup

- > At this point you should be able to communicate with another VM on the same portgroup
 - If not
 - Look at the receive and transmit byte counters in the VMs to see what is going on
 - Look at esxtop, esxcfg-info for any dropped packets on the ports
 - Check firewall settings in the guest



Step 3: Check the physical NIC

Check if the uplinks connected to vSwitch and the portgroup are correct

- > Use VI or `esxcfg-vswitch -l`

Check if the physical NIC is receiving and transmitting data

- > Check the RX and TX counter using `ifconfig`, `esxcfg-info` or `esxtop`
- > Check if Tx/Rx error counters are incrementing
- > Check the VMkernel logs for any errors

Check if the physical NICs are connected to the correct physical switch ports

- > Use CDP or Network Hints

Step 4: Check VLAN Configuration

Check who is tagging and stripping the VLAN IDs?

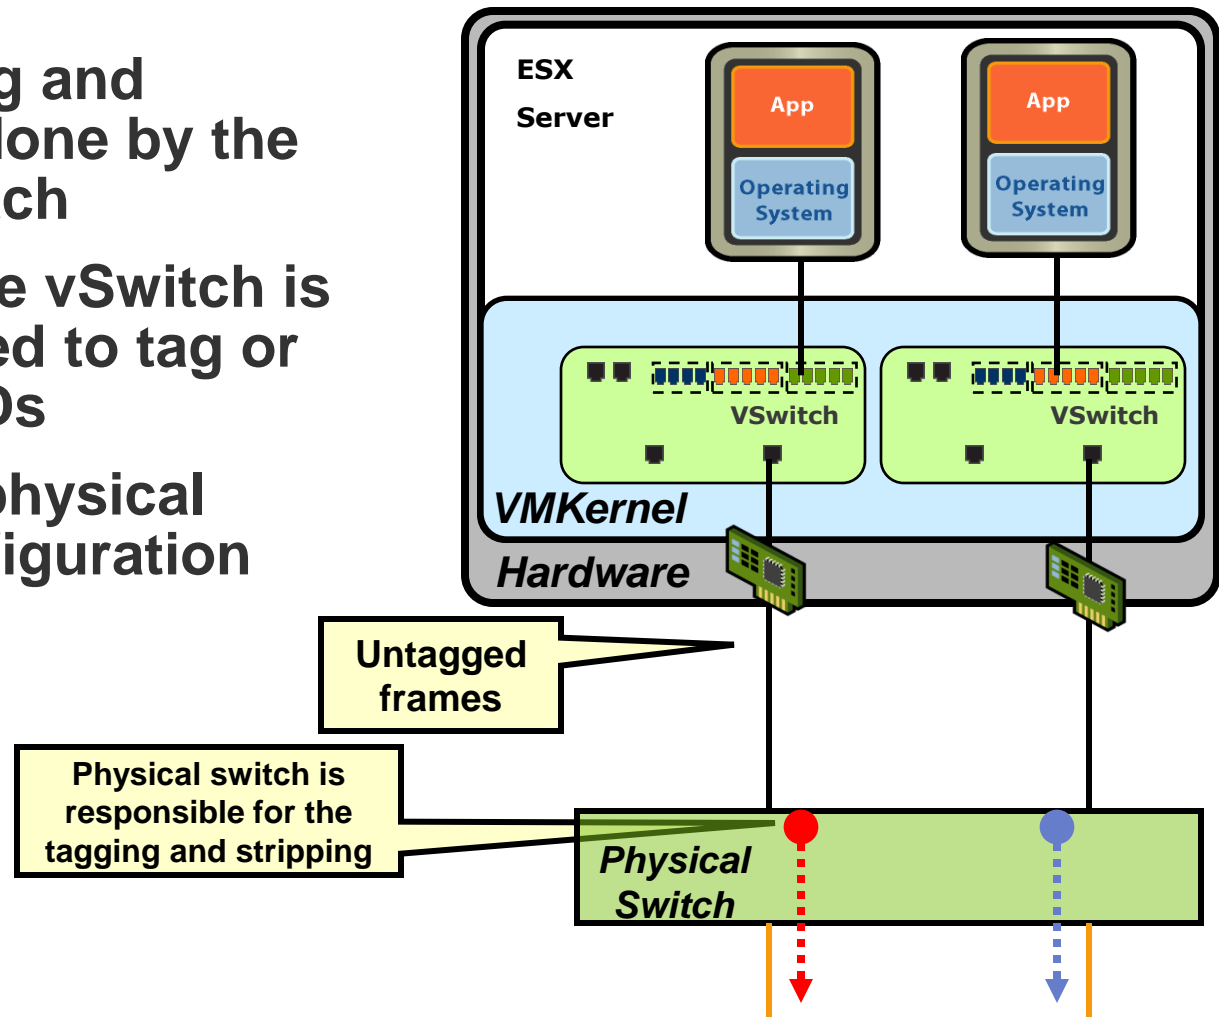
- > External Switch Tagging
 - Only the physical switch tags and strips VLAN IDs
- > Virtual Switch Tagging
 - Only the vSwitch tags and strips VLAN IDs
- > Virtual Guest Tagging
 - Only the guest tags and strips VLAN IDs

Step 4: Check VLAN Configuration: EST

VLAN tagging and stripping is done by the physical switch

Make sure the vSwitch is not configured to tag or strip VLAN IDs

Check your physical network configuration

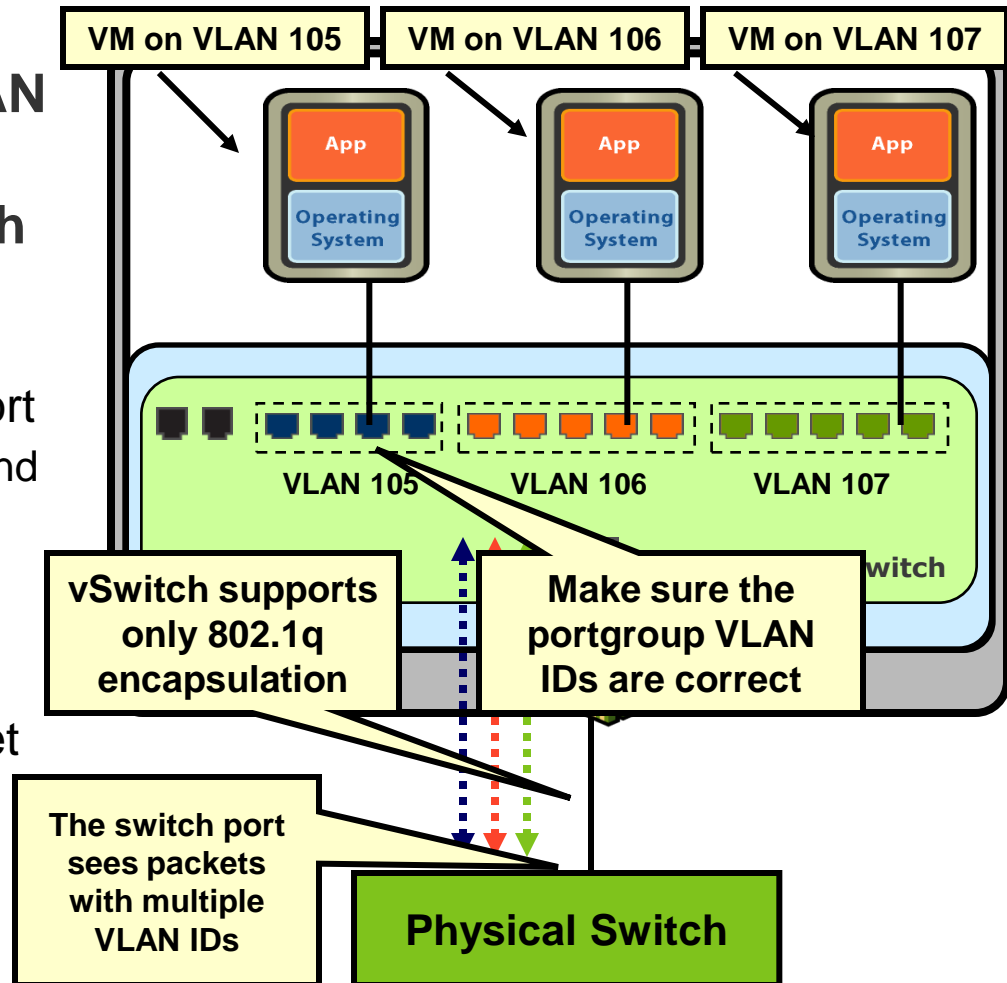


Step 4: Check VLAN Configuration: VST

Check the portgroup VLAN ID

Check the physical switch port configuration

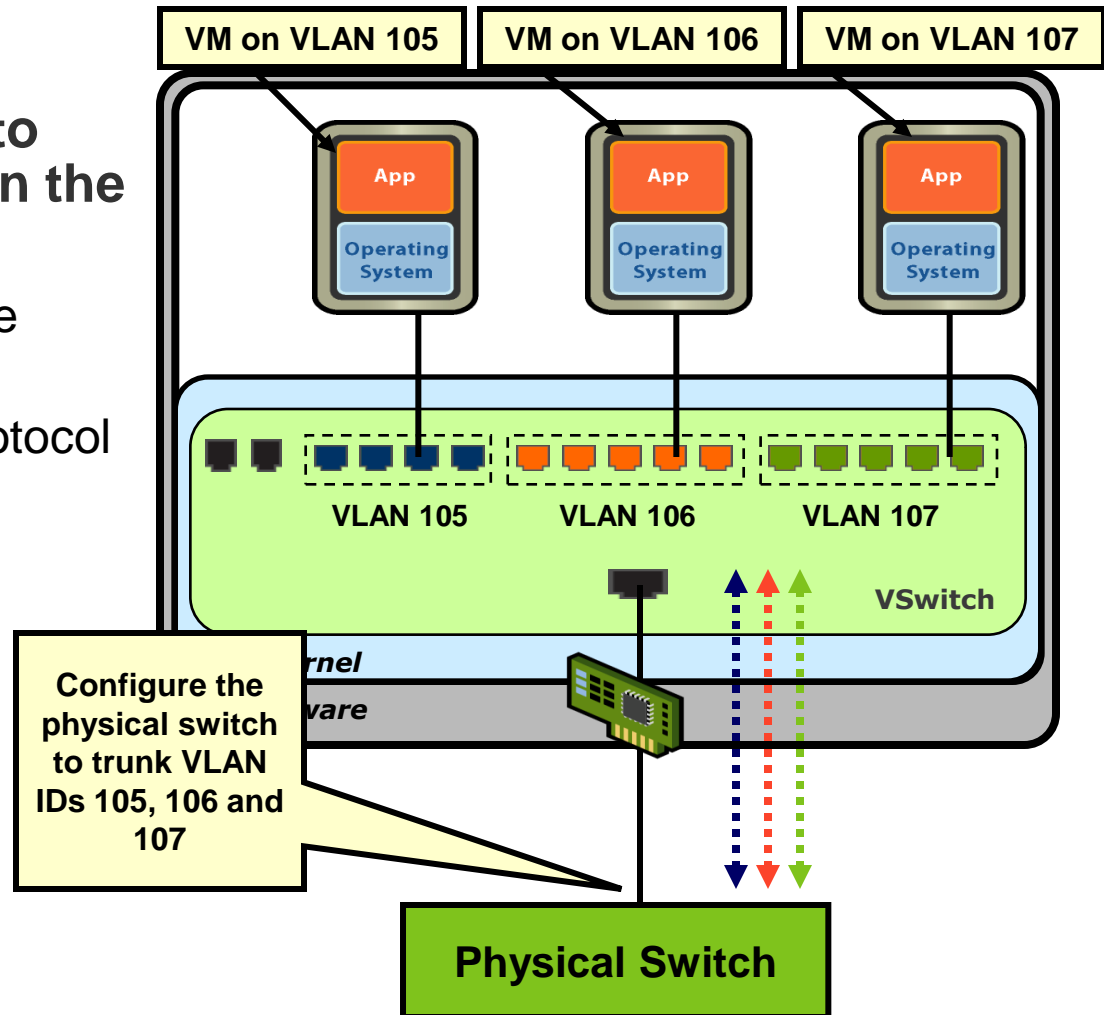
- > Physical switch port should be configured as a trunk port
 - Trunking should be static and unconditional
 - No Dynamic Trunking Protocol (DTP)
- > Physical switch port trunk encapsulation should be set to 802.1q
 - No ISL, LANE etc



Step 4: Check VLAN Configuration: VST

Check if the physical switch is configured to trunk all the VLANs on the vSwitch

- > Manually specify all the VLANs to be trunked
- > No VLAN Trunking Protocol (VTP)

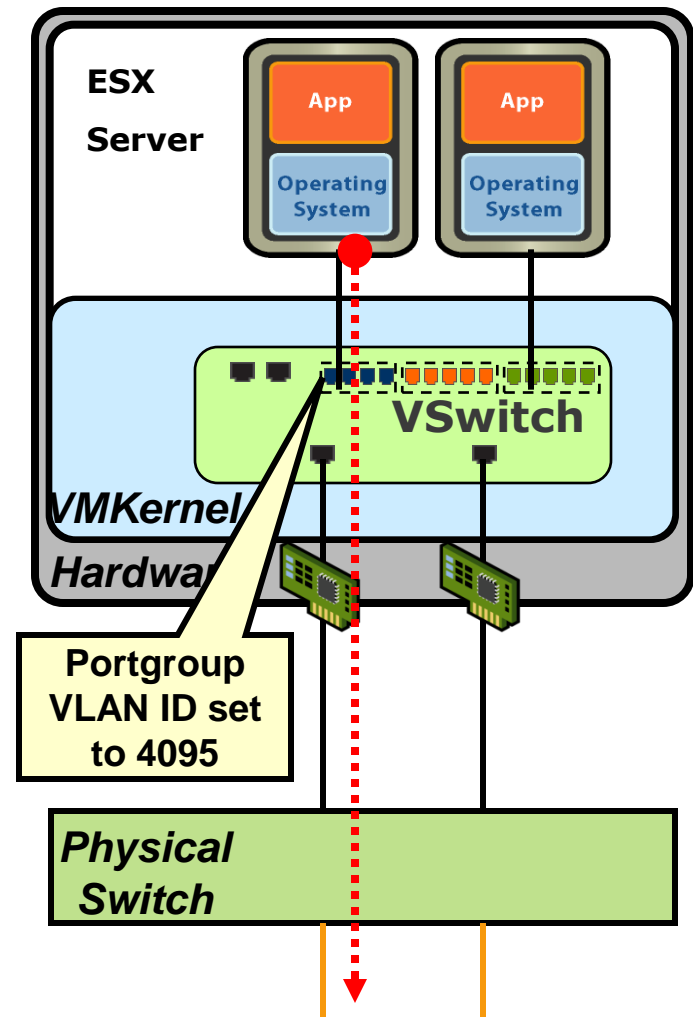


Step 4: Check VLAN Configuration: VGT

Check if the portgroup VLAN Id is set to 4095

Check physical switch configuration

- > Physical switch port should be a statically trunked
- > Physical switch should be configured to expect frames with the specific VLAN IDs on the port
- > Physical switch port trunk encapsulation should be set to 802.1q



Step 4: Check VLAN Configuration: Native VLAN

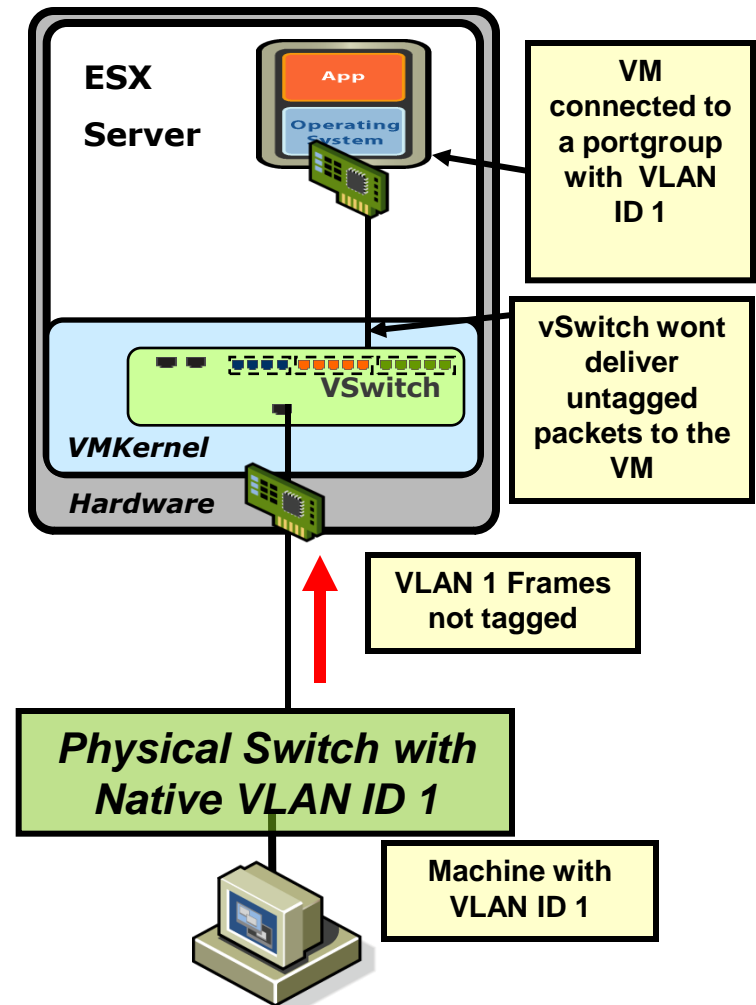
Don't use native VLAN for regular traffic

- > Default native VLAN is often VLAN 1

If you have to use default native VLAN for regular data traffic, do one of the following:

- > Change the native VLAN on the physical switch
- > Force tagging of native VLAN frames

Might need to change native VLAN behavior on all neighboring switches



Problem

Some of the VMs on a vSwitch have network connectivity, other don't

Step 1: Round up the Usual Suspects

Check the vNIC on the VM

- > Check if the vNIC is connected to the correct portgroup
- > Check if VM to VM traffic on the same portgroup works

Check if the physical NIC is connected to the right port/switch


- > Use CDP

Step 2: NIC Teaming

The two VMs could be using different uplinks

NICs in a team should be connected to the same broadcast domain

> Look at the Network Hint



Step 2: NIC Teaming

VI

esxcfg-info

- Search for 'Network Hint' in the output

```
\==+Physical Nic :  
|----Name.....vmnic0  
|----PCI Bus.....1  
|----PCI Slot.....4  
|----PCI function.....0  
|----MAC Address.....00:11:85:13:79:9a  
|----Virtual MAC Address.....00:50:56:53:79:9a  
|----Driver.....tg3  
|----Network Hint.....0 10.17.40.00/255.255.254.00
```

Getting Started

Summary

Virtual Machines

Resource Allocation

Performance

Configuration

Users & Groups

Events

Permissions

Hardware

Health Status

Processors

Memory






Storage

Networking

Storage Adapters

▶ Network Adapters


Network Adapters

Device	Speed	Configured	vSwitch	Observed IP ranges	Wake on LAN Supported
NC7781 Gigabit Server Adapter (PCI-X, 10,100,1000-T)					
 vmnic0	100 Full	Negotiate	vSwitch0	10.17.40.1-10.17.43.254	Yes
 vmnic1	100 Full	Negotiate	None	10.17.40.1-10.17.43.254	Yes
8254NXX Gigabit Ethernet Controller					
 vmnic2	100 Full	Negotiate	None	10.17.40.1-10.17.43.254	Yes
EtherExpress PRO/100 S Server Adapter					
 vmnic3	100 Full	Negotiate	None	10.17.40.1-10.17.43.254	Yes
 vmnic4	100 Full	Negotiate	None	10.17.40.1-10.17.43.254	Yes

Step 3: VLAN Configuration

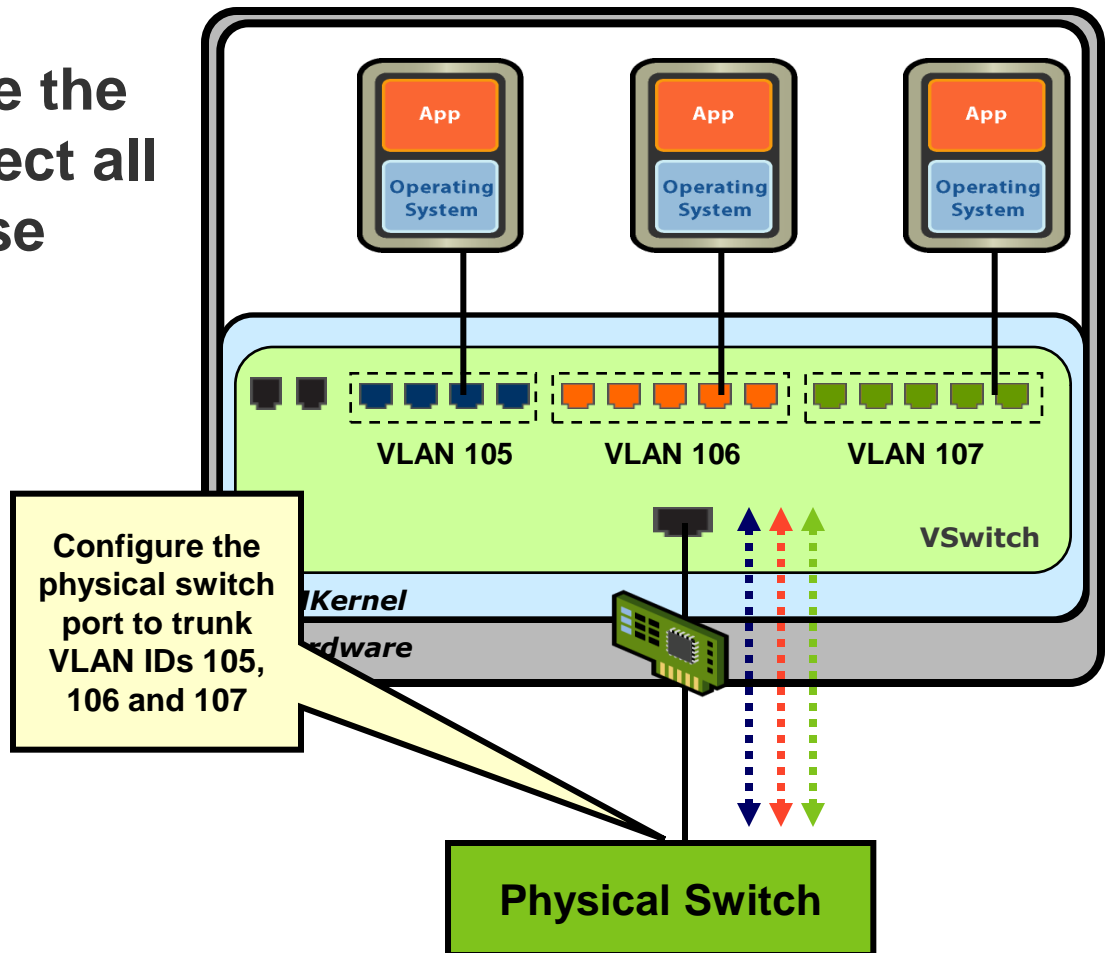
The two VMs could be using different uplinks

VLAN configuration on physical switch ports connected to NICs in a team should be identical



Step 3: VLAN Configuration

Manually configure the switch port to expect all the VLAN IDs in use



Problem

VMs have intermittent network connection

Step 1: Round up the Usual Suspects

Check the vNIC on the VM

- > Check if VM to VM traffic on the same portgroup works without intermittent problems

Check VLAN configuration

- > Identical VLAN configuration on physical switch ports that are in a team

Make sure the NICs in a team are in the same layer 2 broadcast domain

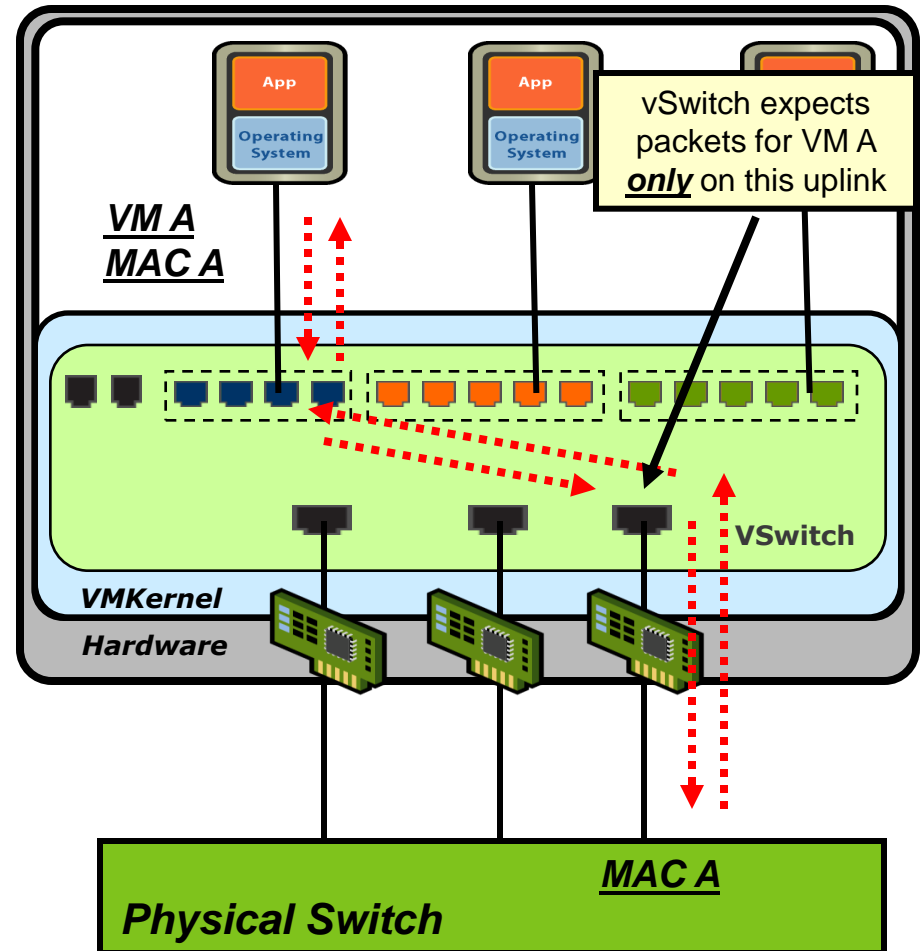
- > Check if the physical NIC is connected to the right port/switch

Step 2: NIC Teaming

Port Id or MAC based
load balancing on ESX



Don't enable Link
Aggregation on the
physical switch



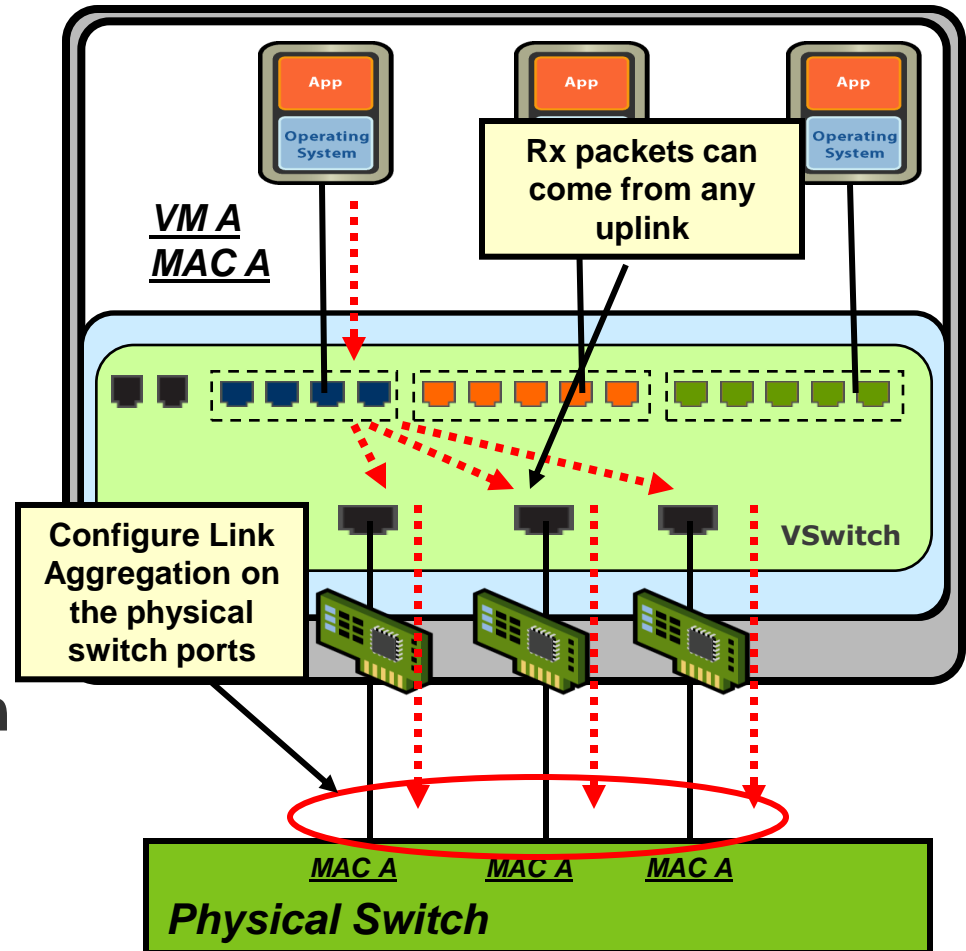
Step 2: NIC Teaming

IP based load balancing
on ESX



Enable Link Aggregation
on the physical switch

- > Static Link Aggregation
 - No LACP or PAgP

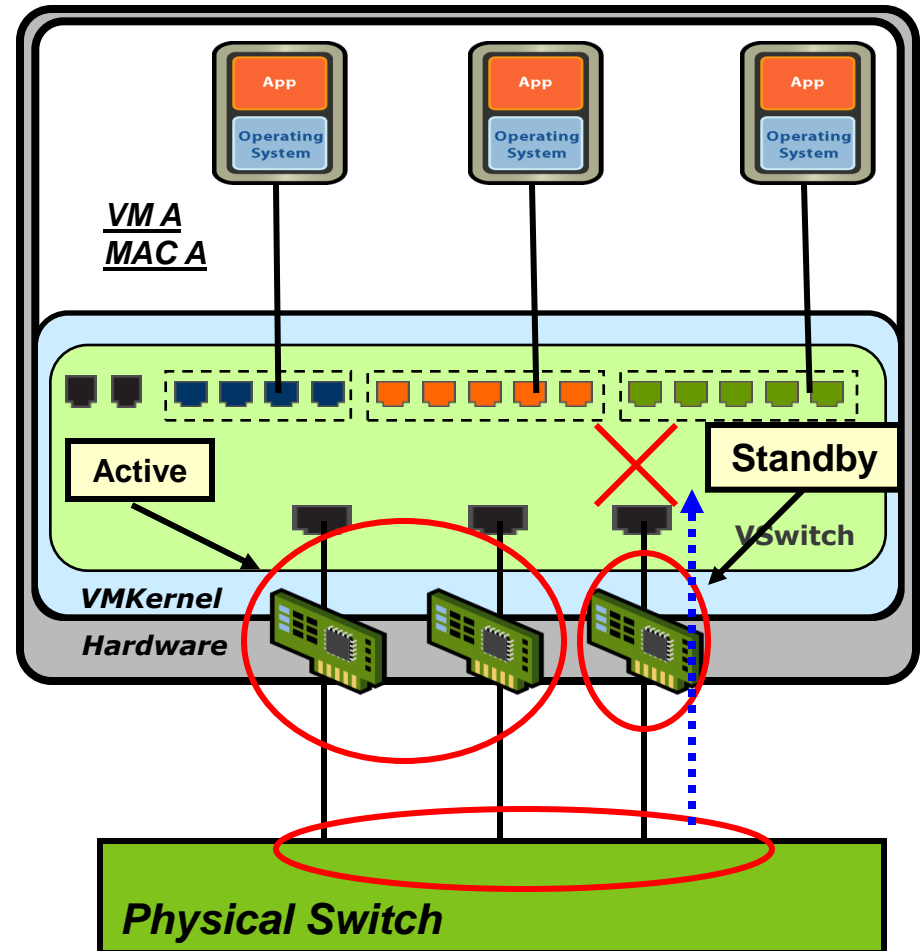


Step 2: NIC Teaming

Active-Standby won't work with IP based load balancing

- Because of the static Link Aggregation the physical switch will want to deliver packets on the standby NIC

Be careful when configuring IP Hash based teaming with other load-balancing configurations on portgroups of the same vSwitch



Step 2: NIC Teaming

Multicast traffic?



Don't use MAC Address Based Load balancing


> Use Port Id based load balancing instead

Step 3: Jumbo Frames

Network devices drop frames larger than MTU

How to detect such problems?

- > Run ping with large packet size, e.g 20000
 - `ping -s <packet size>`
- > Check for packet drops on the vNICs and the physical NICs




Step 3: Jumbo Frames

MTU should be the same end to end

- > Set vNIC MTU in the guest
- > Use `esxcfg-vswitch` to set the MTU of the physical NIC

`esxcfg-vswitch -m <MTU> <vSwitch>`

- Use RCLI for ESXi 3.5



Problem

VMs lose network connectivity upon teaming failover/failback

Step 1: Round up the Usual Suspects

Check physical switch side VLAN configurations

- > Should be identical for all the NICs in a team


Check physical NIC connections

- > NICs in a team must be in the same broadcast domain


Step 2: Spanning Tree Protocol

The switch drops packets on a newly active port till the port is in forwarding state, if STP is enabled

This interferes with failbacks




STP States of a newly active port



The switch is dropping packets on the port till the port is in Forwarding State

Step 2: Spanning Tree Protocol

**Loops are not possible
inside ESX**



Step 2: Spanning Tree Protocol

To avoid the dropped packets, do one of the following

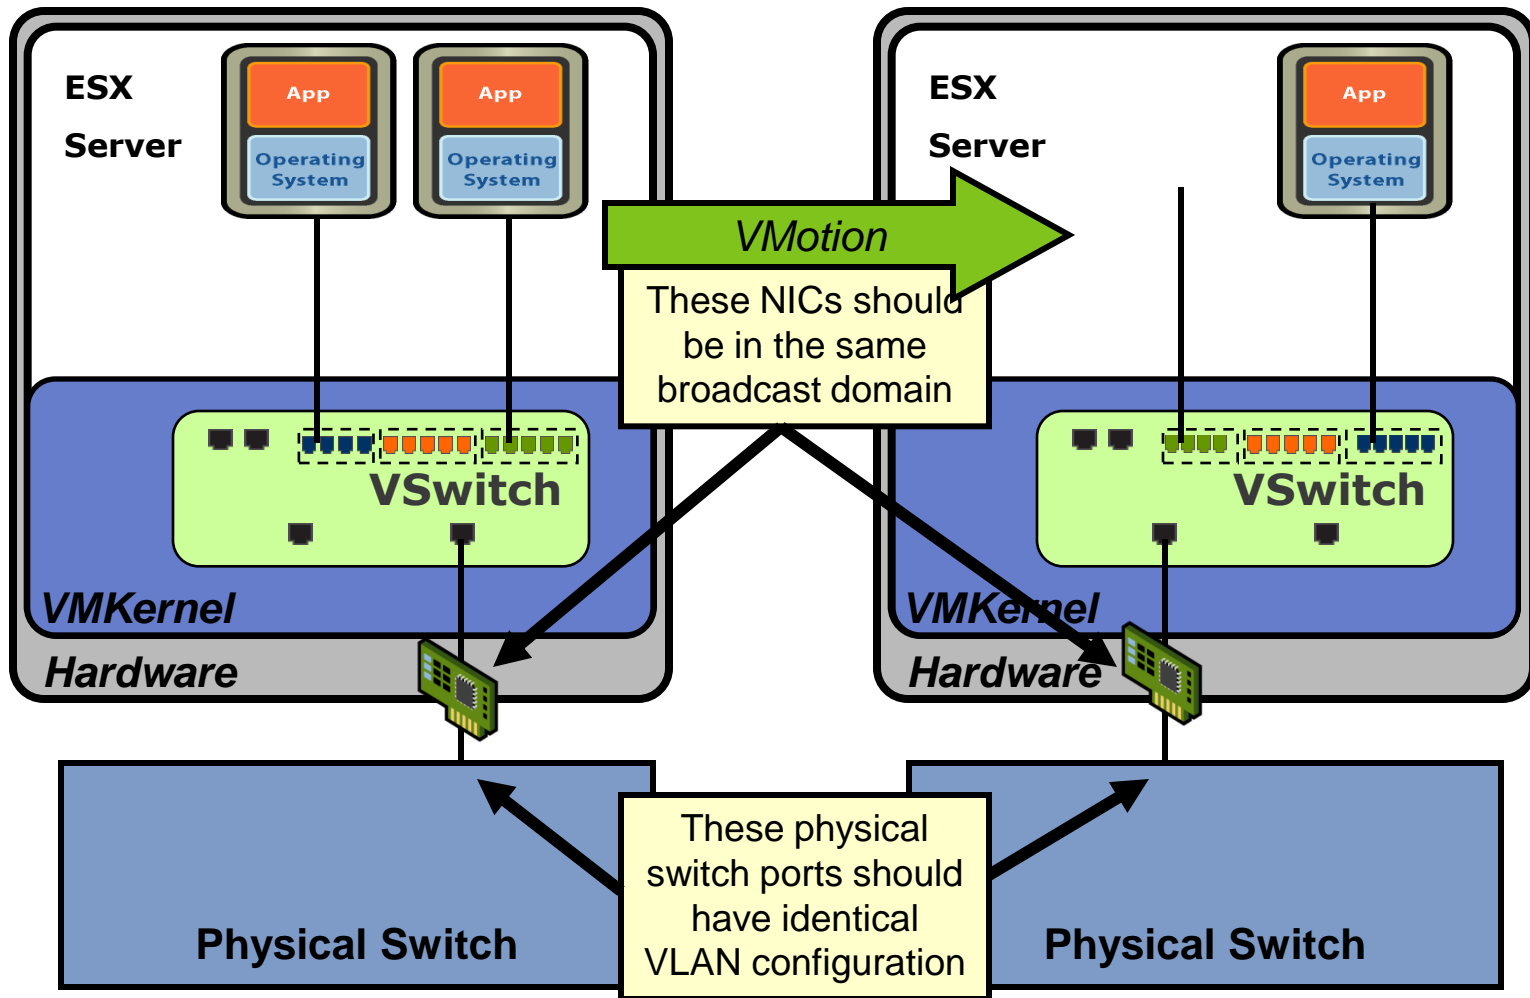
- > Enable PortFast mode for the physical switch ports feeding the ESX Server
- > Configure the physical switch ports feeding the ESX Server as Edge Ports when using Rapid Spanning Tree Protocol
- > Disable STP for the physical switch ports feeding the ESX Server

This is not a recommendation to disable STP in the entire network

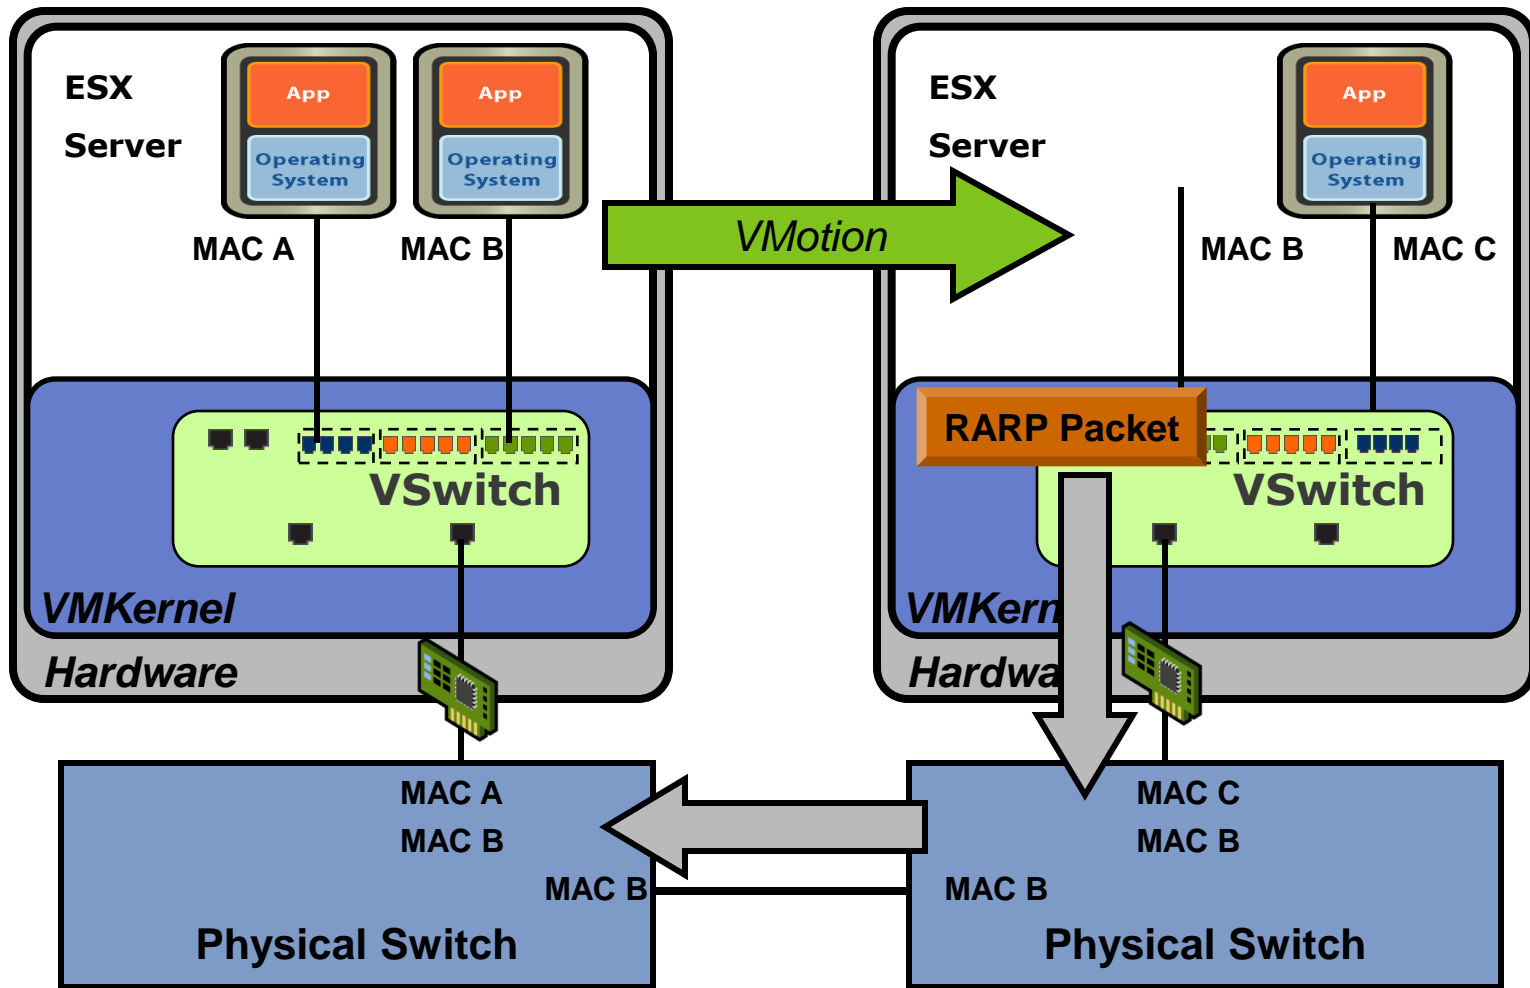
Problem

VMs lose network connectivity after VMotion

Step 1: Basics



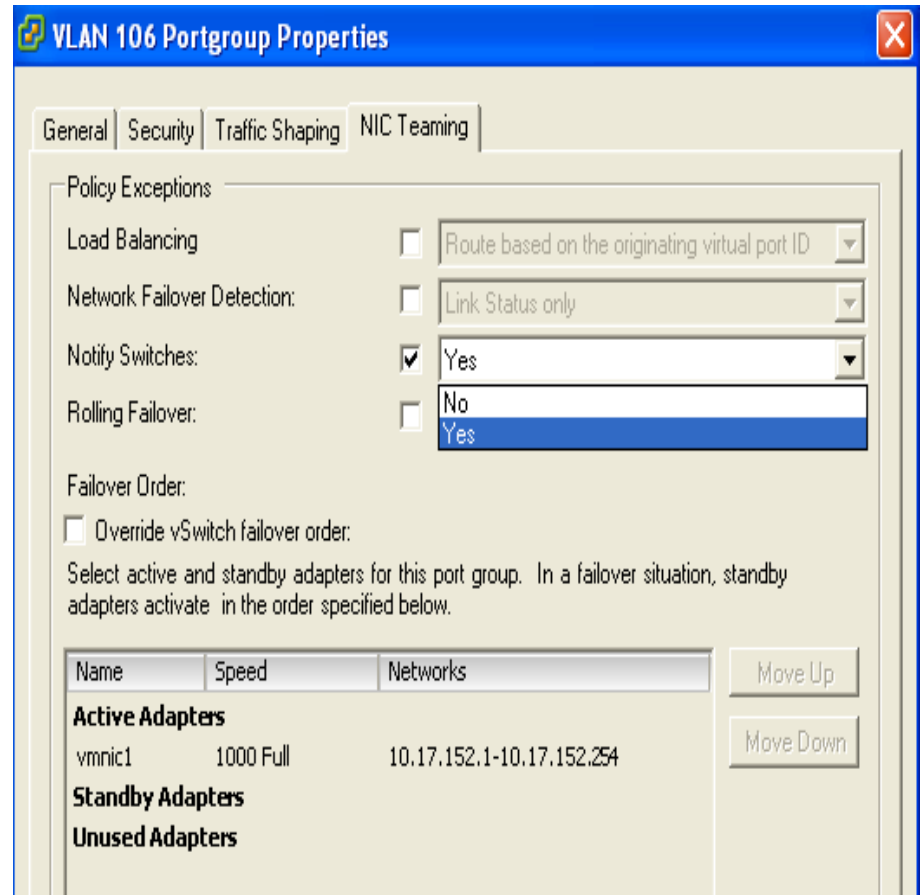
Step 2: Notify Switch



Notify Switch

Notify switch is enabled by default

Settings should reflect application requirements



Problem

Network connectivity is fine on the VMs but my application does not work

Step 1: Check Portgroup Security Policies

Promiscuous Mode

- > If allowed, guest receives all frames on the vSwitch
- > Some applications need promiscuous mode
 - Network sniffers
 - Intrusion detection systems

MAC Address Change

- > If allowed, guest can change its MAC address
 - Implication: Malicious guests can spoof MAC addresses

Forged Transmits

- > If allowed, guest can send packets with different source MAC
 - Implication: Malicious guests can spoof MAC addresses or cause MAC Flooding

Security settings should reflect application requirements

Example: Microsoft Network Load Balancing In Unicast Mode

All cluster hosts are assigned the same MAC address

- > Thus incoming packets are received by all cluster hosts

Uses forged MAC addresses to hide the cluster MAC address from the switch

- > Prevents the switch from learning the cluster's actual MAC address
- > Incoming packets for the cluster are delivered to all switch ports

Portgroup configurations

- > Allow MAC address changes
- > Allow Forged Transmits
- > Do not Notify Switch

KB Article 1556

- > <http://kb.vmware.com/kb/1556>

Recommendation: Use NLB in Multicast Mode



For more information:

VMware Networking Technology

vmware.com/go/networking

VMware Networking Blog

blogs.vmware.com/networking