

Troubleshooting VMware System Problems I

Greg Lindley
Team Lead
IBM VMware Support

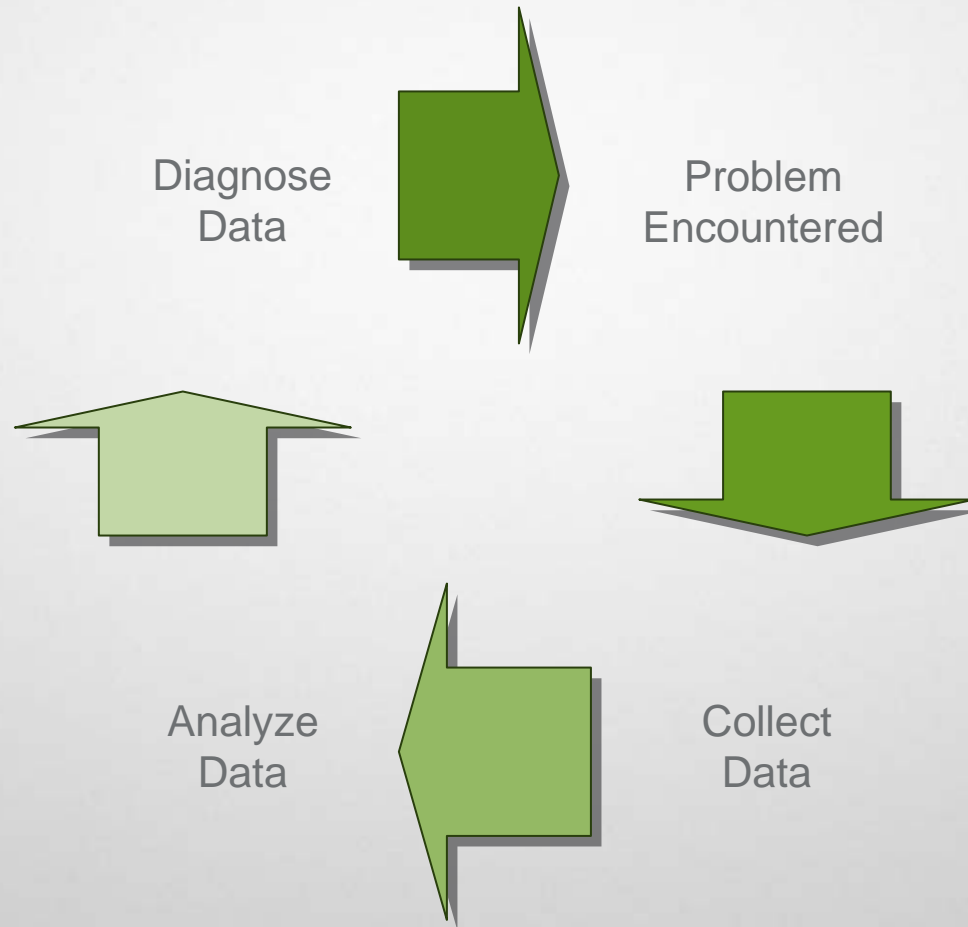
Matrix of Common VMware Problems

	Network	Storage
Physical Machine	Network Problems on the Physical Machine	Storage Problems on the Physical Machine
Virtual Machines	Network Problems on the Virtual Machine	Storage Problems on the Virtual Machine

Source of Other Potential Issues

- Problems with the httpd.vmware server (MUI)
- Root filesystem full
- Intermittent problems (CPU resources, network switch port issues, etc.)
- Hardware problems:
 - Ensure you are at the latest BIOS.
 - Watch for ServeRAID BIOS/driver mismatches
 - Review hardware logs on the system
 - Review hardware logs on the SAN
- PSOD on the Service Console

Steps in Problem Diagnosis



Collect Data

Collecting Diagnostic Data

- Collecting debug information for VMware ESX Server:
 - <http://www.vmware.com/support/kb> (Answer ID 653)
 - Before running the vm-support script, ensure you have the latest one from the link above.
- Collecting debug information for VMware GSX Server:
 - <http://www.vmware.com/support/kb> (Answer ID 1345)
- Collecting debug information for VMware VirtualCenter:
 - <http://www.vmware.com/support/kb> (Answer ID 1365)

Collect Data

Collecting Diagnostic Data

- Collect SAN log information
 - Examples:
 - FASStT MEL logs
 - FASStT profile information
- Collect network switch port information
 - Have network support check for port failures, checksums errors, etc., during the time of the failure
- It is best to collect diagnostic data at the time of the failure. Otherwise critical information may be lost

Analyze Data

Layout of Virtual Machine Support

- Extract the virtual machine support with the following command:
 - `tar zxvf esx-XXXX-XX-XX.XXXX.tgz`
- Directories captured with the current script:
 - `etc/` - various config information for VMware
 - `proc/` - snapshot of the system state
 - `root/` - only if the machine had a PSOD or if `.vmx` files are located there
 - `tmp/` - config information
 - `usr/` - `httpd.vmware` apache config information
 - `var/` - various system logs
 - and possibly `'home'` or `'vpx'` depending on the location of the `.vmx` config files (see `etc/vmware/vm-list`)
- Always start in the `tmp` directory to get an overall view of the system setup

Analyze Data

Determining What PCI Cards Are in the System

```
# /sbin/lspci
```

```
00:00.0 Host bridge: ServerWorks: Unknown device 0014 (rev 33)
00:00.1 Host bridge: ServerWorks: Unknown device 0014
00:00.2 Host bridge: ServerWorks: Unknown device 0014
00:01.0 VGA compatible controller: ATI Technologies Inc Radeon VE QY
00:0f.0 ISA bridge: ServerWorks: Unknown device 0203 (rev a0)
00:0f.1 IDE interface: ServerWorks: Unknown device 0213 (rev a0)
00:0f.2 USB Controller: ServerWorks: Unknown device 0221 (rev 05)
00:0f.3 Host bridge: ServerWorks: Unknown device 0227
00:10.0 Host bridge: ServerWorks: Unknown device 0101 (rev 05)
00:10.2 Host bridge: ServerWorks: Unknown device 0101 (rev 05)
01:01.0 Ethernet controller: Intel Corporation 8254NXX Gigabit Ethernet Controller (rev 04)
01:02.0 Ethernet controller: Intel Corporation 8254NXX Gigabit Ethernet Controller (rev 04)
02:01.0 Ethernet controller: Intel Corporation 8254NXX Gigabit Ethernet Controller (rev 03)
02:01.1 Ethernet controller: Intel Corporation 8254NXX Gigabit Ethernet Controller (rev 03)
02:03.0 Fiber Channel: QLogic Corp QLA231x/2340 (rev 02)
02:03.1 Fiber Channel: QLogic Corp QLA231x/2340 (rev 02)
```

Also in `tmp/lspci1.*.txt`

First field is Bus:Slot:Function

Lines with same Bus/Slot, but different Function are dual/quad port cards.

Analyze Data

Determining How Devices Are Allocated

```
# /usr/sbin/vmkchdev -L
000:00.0 1166:0014 0000:0000 console
PCI device 1166:0014 (ServerWorks)
000:00.1 1166:0014 0000:0000 console
PCI device 1166:0014 (ServerWorks)
000:00.2 1166:0014 0000:0000 console
PCI device 1166:0014 (ServerWorks)
...
001:01.0 8086:1028 8086:34b1 vmkernel vmnic3
PCI device 8086:1028 (Intel Corporation)
001:02.0 8086:1028 8086:34b1 console vmnic0
PCI device 8086:1028 (Intel Corporation)
```

Also in tmp/vmkchdev.*.txt

Devices that say “console” are allocated to the console.

Devices that say vmkernel are allocated to the vmkernel.

Devices that say console and have a vm device designation are shared.

Same information is also in etc/vmware/hwconfig (i.e. host, virtual machine, shared).

Analyze Data

Determine What Drivers are Loaded in the Service Console

```
# /sbin/lsmmod
```

Module		Size	Used by Tainted: PF
vmxnet_console		13212	1
Vmnixmod	177056	3 [vmxnet_console]	
E1000		68456	0 (unused)
usb-storage		20028	0
Mousedev	3936	0 (unused)	
Keybdev		1696	0 (unused)
Hid		17728	0 (unused)
Input		3488	0 [mousedev keybdev hid]
usb-ohci	17600	0 (unused)	
Usbcore		50112	1 [usb-storage hid usb-ohci]

Also in tmp/modules.*.txt

Analyze Data

Determine What Drivers Are Loaded by the vmkernel

```
# /usr/sbin/vmkload_mod -l
```

Name	Loaded	R/O Addr	Length	R/W Addr	Length	ID
Vmklinux		0x4de000	0xf000	0x12438f8 0x53000	1	Yes
Nfshaper		0x4ed000	0x1000	0x129b160 0x1000	2	Yes
E1000		0x4ee000	0xf000	0x129c168 0x6000	3	Yes
qla2300_604		0x4fd000	0x19000	0x12fe008 0x22000	4	Yes
Bond		0x516000	0x2000	0x1574b80 0x2000	5	Yes

Also in tmp/vmkmod.*.txt

Not all drivers can be unloaded/loaded successfully with the vmkload_mod command.

Analyze Data

Determine How the SAN Is Presented

```
# /usr/sbin/vmkmultipath -q
```

Disk and multipath information follows:

Disk vmhba0:0:0 (225,278 MB) has 4 paths. Policy is mru.

vmhba0:0:0 on (active, preferred)

vmhba0:1:0 on

vmhba1:0:0 on

vmhba1:1:0 on

Disk is named after the first channel it is seen on.

If active != preferred, there could be SAN issues to investigate.

Check device numeration in proc/scsi.

Also in tmp/vmkmultipath.*.txt

Analyze Data

Determine How the Network Is Set Up

```
# cat etc/vmware/netmap.conf
network0.name = "External Linux Net"
network0.device = " Public Web access.20"
network1.name = "Internal Windows Net"
network1.device = "vmnet_0"
network2.name = "Public Web access"
network2.device = "bond0"
```

Real devices are listed as vmnic? and bond?.
Vmnet devices are pseudo devices.

Device designations ending in .# (e.g. bond0.610)
contain the VLAN IDs on the end.

Analyze Data

Examine etc/vmware/hwconfig for More
Network Information

- Sample text from file:
 - nicteam.vmnic0.team = "bond0"
 - nicteam.vmnic1.team = "bond0"

Analyze Data

System Network Troubleshooting

```
# cat proc/vmware/net/vmnic0/config
VLanHwTxAccel Yes
VLanHwRxAccel Yes
VLanSwTagging Yes
PromiscuousAllowed No
InterruptClustering No
Link state: Up Speed: 1000 Mbps, full duplex
Queue: Running PCI (bus:slot.func): 1:0.0
Minimum Capabilities 0x0
Device Capabilities 0x74b
Maximum Capabilities 0x76b
NICTeamingMaster: bond0
TeamFailoverBeacon: Off
Interrupt vector 0x69
DebugSocket Closed
```

The config file in the bond directories will tell you if the bond is setup for load balancing, out-mac, or out-ip.

Diagnose Data

System Network Troubleshooting

Connectivity issues with switches (check logs in var/log):

```
Feb 2 12:48:23 SYNVM7 kernel: bcm5700: eth0 NIC Link is DOWN
```

```
Feb 2 12:48:23 SYNVM7 kernel: bcm5700: eth0 NIC Link is UP, 1000 Mbps full duplex
```

You may need to hard code the speed/duplex if the card is having trouble negotiating with the switch.

Serial over LAN/ipmi on HS20?:

<http://www.vmware.com/support/kb> (Answer ID 1627)

Net.ZeroSpeedLinkDown set?:

<http://www.vmware.com/support/kb> (Answer ID 1609)

Diagnose Data

System Network Troubleshooting

VLAN issues:

<http://www.vmware.com/support/kb> (Answer ID 1655)

<http://www.vmware.com/support/kb> (Answer ID 1266)

Try booting up Linux to see if you encounter the same problem.
Verify connectivity to all paths with ping.

Ensure basic setup works before adding complexity. Start with a vswitch or vmnic setup, add bonding, then VLAN. Test connectivity each step of the way.

Diagnose Data

Virtual Machine Network Troubleshooting

- Connectivity issues:
 - Verify that it is not a system problem
 - Check that ping works on the network
 - Try a different driver (vlance/vmxnet)
 - Reinstall the vmware tools
 - Check route, gateway, etc., within the guest

Diagnose Data

System Storage Troubleshooting

- If you are using clustering, ensure the system is set up correctly:
 - http://www.vmware.com/support/esx25/doc/admin/esx25admin_cluster_setup_esx.html
- Make sure the environment is supported:
 - http://www.vmware.com/pdf/esx_SAN_guide.pdf
- Active/Active SAN devices should have a fixed path policy. Active/Passive should be set to MRU

Diagnose Data

System Storage Troubleshooting

- Gather zoning information
 - Ensure all HBAs can see all LUNs
- For FASStT, make sure the host type is set correctly (LNXCL) to prevent LUN “ping-ponging” – LNXCL turns off ADT/AVT
- Make sure the HBAs are dedicated to the virtual machines and not shared with the console (except in BFS situations)

Diagnose Data

System Storage Troubleshooting

- Review var/log/vmkernel for SCSI sense messages:
 - Mar 23 18:11:11 nbxva016 vmkernel: 7:02:38:17.586 cpu5:156)
WARNING: SCSI: 4226: vmhba2:3:11:1 status = 8/2 0x0 0x0 0x0
- Review var/log/vmkernel for “bad” codes:
 - Nov 30 14:07:30 vms-2 vmkernel: 21:22:44:32.195 cpu3:133)
WARNING: SCSI: 7180: returns error: "I/O error". Code: 0xbad000a
 - Also note, that some of these return codes are represented in decimal form:
**Nov 30 14:02:04 vms-2 vmkernel: 21:22:39:06.943 cpu0:132) WARNING:
SCSI: 5303: Failed for vmhba3:1:1 status = 195887114**
 - In the example above, 195887114 = 0xbad000a
- Escalate to your support organization to decode the above

Diagnose Data

Virtual Machine Storage Troubleshooting

- Verify that it is not a system problem
- Make sure the file exists and/or the virtual scsi adapter is pointed at the right device:
 - Check the vmhba.*.txt file in tmp
- Adjust the timeout for Windows if you encounter bluescreens:
 - <http://www.vmware.com/support/kb> (Answer ID 1014)
- Virtual machine panics with bugNr=27436
 - Switch to the Isilogic controller until ESX Server V2.5.2

Diagnose Data

Virtual Machine Storage Troubleshooting

Access to storage:

Make sure the device isn't in use by another machine:

```
Feb 02 11:39:07: vmx| DiskVmnixSetupSCSIDevices: failed to get handle for SCSI Device 0
Feb 02 11:39:07: vmx| Msg_Post: Error
Feb 02 11:39:07: vmx| [msg.diskVmnix.scsiopenfailed] Unable to open scsi target VMFS-
SAN1:mywindowsmachine.vmdk: Device or resource busy(2)
Feb 02 11:39:07: vmx| [msg.vmxlsilogic.poweronFailed]
Feb 02 11:39:07: vmx| Failed to configure scsi0.
Feb 02 11:39:07: vmx| -----
Feb 02 11:39:07: vmx| POST(no connection): Unable to open scsi target VMFS-SAN1:
mywindowsmachine.vmdk: Device or resource busy(2)
Feb 02 11:39:07: vmx|
Feb 02 11:39:07: vmx| Failed to configure scsi0.
Feb 02 11:39:07: vmx|
Feb 02 11:39:07: vmx| Module VmxLSILogic power on failed.
```

Use `ps -efwww` or look at `tmp/ps.*.txt` output:

```
root 2751 0.0 1.7 404252 6612 ? S< 12:29 0:00 vmware-mks -A 11 -D 13 -S -L /tmp/vmware-root-
2750.log -P 2750 -g -@ vm=374709cf368cf239; gui=false; vmdbMemMapHandle=0x4;
vmdbMemMapSize=0x400000; useSELinux=false -C
/home/vmware/mywindowsmachine/mywindowsmachine.vmx
```

Fixing Problems With the MUI

- Restart the MUI service:
 - # service httpd.vmware restart
 - Shutting down http.vmware: [OK]
 - Starting httpd.vmware: [OK]
- Reinstall the MUI:
 - # rpm -e VMware-mui-2.1.2-9638
 - # rpm -i VMware-mui-2.1.2-9638.rpm
- Investigate the logs in /var/log/vmware-mui
- Can you telnet to the port (80)?

Resolving Other Issues

- Root filesystem full
 - Clean out old logs in /var/log
 - Configure logrotate
 - Increase frequency
 - Include additional logs (MUI error_log)
 - /etc/vmware/config may become NULL
- Several BIOS revisions have contributed to random PSODs. Ensure you are at the latest level
- ServeRAID driver is different between VMware releases. Ensure that a flash of the ServeRAID firmware is part of your upgrade plans

Resolving Other Issues, cont.

- Intermittent problems:
 - Network connectivity:
 - Check switch logs
 - Ensure you are not running out of CPU resources in the console OS
 - SAN connectivity:
 - Check switch and SAN device logs
 - Make sure SAN device is connected per VMware recommendations
- PSOD on the Service Console
 - Always get a screenshot in case the dump doesn't work
 - Escalate to appropriate Service organization to resolve

Additional References

- How to troubleshoot VMware system problems
 - <http://www-128.ibm.com/developerworks/linux/library/l-vmware/>
- Implementing VMware ESX Server 2.1 with IBM TotalStorage FAStT
 - <http://www.redbooks.ibm.com/abstracts/sg246434.html?Open>
- IBM 2105 (Shark):
 - <http://publibfp.boulder.ibm.com/epubs/pdf/f2bhs03.pdf>
- IBM 2145 (SVC):
 - <http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001144&rs=591>
- EMC Clariion
 - <http://powerlink.emc.com>
 - Host Connectivity Guide for VMware ESX Server
- HP MSA/EVA
 - <http://h71019.www7.hp.com/ActiveAnswers/cache/71085-0-0-225-121.html>
 - Configuring redundant HP StorageWorks SAN infrastructure with VMware ESX Server 2.1
- <http://www.vmware.com/kb>
- <http://www.vmware.com/community>
- <http://www.vmware.com/support/pubs>

vmworld2005

virtualize^{now}

las vegas • october 18-20, 2005