

<http://blog.michaelfmcnamara.com/2009/09/vsphere-scsi-reserv-co/>

We had our first issue today with our recent VMware vSphere 4 installation. We're currently up to about 30 virtual machines spread across five BL460c (36GB) blades in an HP 7000 Enclosure. The problem started with a few virtual machines just going south, like they had lost their mind. It was discovered that all the virtual machines that were affected were on the same datastore (LUN). One of the engineers put the ESX host that was running those VMs into maintenance mode and rebooted it. After the reboot the ESX host was unable to mount the datastore. Everything seemed fine from a SAN standpoint and the Fiber Channel switches were working fine. A quick look at /var/log/vmkwarning on the ESX host revealed the following messages;

```
Sep  1 13:04:35 mdcc01h10r242 vmkernel: 0:00:26:02.384
cpu4:4119)WARNING: ScsiDeviceIO: 1374: I/O failed due to too many
reservation conflicts. naa.600508b4000547cc0000b00001540000 (920 0 3)
Sep  1 13:04:40 mdcc01h10r242 vmkernel: 0:00:26:07.400
cpu6:4119)WARNING: ScsiDeviceIO: 1374: I/O failed due to too many
reservation conflicts. naa.600508b4000547cc0000b00001540000 (920 0 3)
Sep  1 13:04:40 mdcc01h10r242 vmkernel: 0:00:26:07.400
cpu6:4119)WARNING: Partition: 705: Partition table read from device
naa.600508b4000547cc0000b00001540000 failed: SCSI reservation conflict
```

A quick examination of the other ESX hosts revealed the following;

```
Sep  1 13:04:26 mdcc01h09r242 vmkernel: 21:22:13:25.727
cpu10:4124)WARNING: FS3: 6509: Reservation error: SCSI reservation
conflict
Sep  1 13:04:31 mdcc01h09r242 vmkernel: 21:22:13:30.715
cpu12:4124)WARNING: FS3: 6509: Reservation error: SCSI reservation
conflict
Sep  1 13:04:36 mdcc01h09r242 vmkernel: 21:22:13:35.761
cpu9:4124)WARNING: FS3: 6509: Reservation error: SCSI reservation
conflict
```

We had a SCSI reservation issue that was locking out the LUN from any of the ESX hosts. The immediate suspect was the VCB host as it was the only other host that was being presented the same datastores (LUNs) as the ESX hosts from the SAN (HP EVA 6000).

We rebooted the VCB host and then issued the following command to reset the LUN from one of the ESX hosts;

```
vmkfstools -L lunreset
/vmfs/devices/disks/naa.600508b4000547cc0000b00001540000
```

After issuing the LUN reset we observed the following message in the log;

```
Sep  1 13:04:40 mdcc01h10r242 vmkernel: 0:00:26:07.400
cpu9:4209)WARNING: NMP: nmp_DeviceTaskMgmt: Attempt to issue lun reset
on device naa.600508b4000547cc0000b00001540000. This will clear any
SCSI-2 reservations on the device.
```

The ESX hosts were almost immediately able to see the datastore and the problem was resolved.

We believe the problem occurred when the VCB host tried to backup multiple virtual machines from the same datastore (LUN) at the same time. This created an issue when the VCB host locked the LUN for too long causing the SCSI queue to fill-up on the ESX hosts. This is new to us and to me so we're still trying to figure it out.

Cheers!

References;

<http://kb.vmware.com/kb/1009899>

http://www.vmware.com/files/pdf/vcb_best_practices.pdf

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayK&externalId=1003742