

problems after SAN failure. ...



AllBlack 397 posts since

Nov 5, 2007 Hi everyone.

On friday we had a SAN failure. Logged a job with EMC and they had no clue as nothing was obvious.

They did notice that both storage processors had a panic at the same time. The EMC development engineers are looking into it. Needless to say that pretty much everything turned to custard. A lot of VMs are unhappy and pretty much needed a cold reboot.

On one of our hosts I have lost one of the LUNs. On the SAN there are two LUNs available. ESX seems to detect the same LUN twice and obviously I have issues with my paths.

Some output that is related

1. esxcfg-mpath -l

Disk vmhba32:2:0 (0MB) has 2 paths and policy of Most Recently Used

iScsi sw iqn.1998-01.com.vmware:localhost-3055e03e<->iqn.1992-04.com.emc:cx.ck200064601253.a2
vmhba32:2:0 On active preferred

iScsi sw iqn.1998-01.com.vmware:localhost-3055e03e<->iqn.1992-04.com.emc:cx.ck200064601253.b2
vmhba32:6:0 On

Disk vmhba0:0:0 /dev/cciss/c0d0 (69973MB) has 1 paths and policy of Fixed

Local 6:0:0 vmhba0:0:0 On active preferred

Disk vmhba32:3:1 /dev/sdb (512000MB) has 2 paths and policy of Most Recently Used

iScsi sw iqn.1998-01.com.vmware:localhost-3055e03e<->iqn.1992-04.com.emc:cx.ck200064601253.a3
vmhba32:3:1 Standby preferred

iScsi sw iqn.1998-01.com.vmware:localhost-3055e03e<->iqn.1992-04.com.emc:cx.ck200064601253.b3
vmhba32:7:1 On active

Disk vmhba32:3:3 (512000MB) has 2 paths and policy of Most Recently Used

iScsi sw iqn.1998-01.com.vmware:localhost-3055e03e<->iqn.1992-04.com.emc:cx.ck200064601253.a3
vmhba32:3:3 Dead preferred

iScsi sw iqn.1998-01.com.vmware:localhost-3055e03e<->iqn.1992-04.com.emc:cx.ck200064601253.b3
vmhba32:7:3 Dead

Disk vmhba32:3:0 /dev/sda (512000MB) has 2 paths and policy of Most Recently Used

iScsi sw iqn.1998-01.com.vmware:localhost-3055e03e<->iqn.1992-04.com.emc:cx.ck200064601253.a3
vmhba32:3:0 Standby preferred

iScsi sw iqn.1998-01.com.vmware:localhost-3055e03e<->iqn.1992-04.com.emc:cx.ck200064601253.b3
vmhba32:7:0 On active

1. esxcfg-vmhbadevs

vmhba0:0:0 /dev/cciss/c0d0

problems after SAN failure. ...

vmhba32:3:0 /dev/sda

vmhba32:3:1 /dev/sdb

These two seem to be the same physical LUN though.

Some log data

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.003 cpu5:1043)WARNING: SCSI: 4541: Delaying failover to path vmhba32:7:3

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.004 cpu1:1025)SCSI: 5270: vml.020003000060060160a2a01a007eea3c745c6edd11

524149442035: Cmd failed. Blocking device during path failover.

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.006 cpu2:1058)SCSI: 2741: Could not locate path to peer SP for CX SP B p

ath vmhba32:7:3.

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.006 cpu2:1058)SCSI: 2741: Could not locate path to peer SP for CX SP B p

ath vmhba32:7:3.

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.006 cpu2:1058)SCSI: 2308: Unmapped LUN state for DGC path vmhba32:7:3

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.007 cpu2:1058)SCSI: 2308: Unmapped LUN state for DGC path vmhba32:3:3

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.007 cpu2:1058)WARNING: SCSI: 4559: Manual switchover to path vmhba32:7:3

begins.

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.007 cpu2:1058)SCSI: 2308: Unmapped LUN state for DGC path vmhba32:7:3

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.007 cpu2:1058)WARNING: SCSI: 3743: Could not switchover to vmhba32:7:3.

Check Unit Ready Command returned an error instead of NOT READY for standby controller .

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.007 cpu2:1058)WARNING: SCSI: 4619: Manual switchover to vmhba32:7:3 comp

leted unsuccessfully.

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.010 cpu5:1057)SCSI: 2741: Could not locate path to peer SP for CX SP B p

ath vmhba32:7:3.

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.010 cpu5:1057)SCSI: 2741: Could not locate path to peer SP for CX SP B p

ath vmhba32:7:3.

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.010 cpu5:1057)SCSI: 2308: Unmapped LUN state for DGC path vmhba32:7:3

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.010 cpu5:1057)SCSI: 2308: Unmapped LUN state for DGC path vmhba32:3:3

problems after SAN failure. ...

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.010 cpu5:1057)WARNING: SCSI: 4559: Manual switchover to path vmhba32:3:3

begins.

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.010 cpu5:1057)jSCSI: session 0xba402c0 eh_device_reset at 1589761539 for

command 0x6636888 to (0 0 3 3), cdb 0x0

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.010 cpu2:1081)jSCSI: session 0xba402c0 requested target reset for (0 0 3

*), warm reset itt 25080319 at 1589761539

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.016 cpu6:1082)jSCSI: session 0xba402c0 warm target reset success for mgm

t 25080319 at 1589761539

Feb 21 14:01:03 tur-esx-dev1 vmkernel: 184:00:00:15.017 cpu2:1081)jSCSI: session 0xba402c0 (0 0 3 *) finished reset at 15897

Some dmesg output

VMWARE: Device that would have been attached as scsi disk sda at scsi1, channel 0, id 2, lun 0

Has not been attached because this path is not active.

key = 0x2, asc = 0x4, ascq = 0x1

VMWARE: Device that would have been attached as scsi disk sda at scsi1, channel 0, id 2, lun 0

Has not been attached because it is a duplicate path or on a passive path

I have never dealt with such an issue so any pointers would be appreciated.

Cheers Tags: unmapped, lun, sp



Lightbulb 1,391 posts since

Aug 15, 2008 1. **Re: problems after SAN failure.** Feb 23, 2009 3:26 AM

I take that your other hosts (Is this a cluster?) are fine and the VMs are running on those hosts, is this correct?

You could try step 3 from the following document that deals with cleaning up iSCSI config on a ESX system

http://apps.sourceforge.net/mediawiki/iscsitarget/index.php?title=The_case_of_stale_iSCSI_LUNs

If your VMs are safely on other hosts you may want to evict this host and reinstall ESX and add host back to cluster. Kind of a cop out but may be the best use of your time. Of course these suggestions are predicated on your VMs running on another host that is not having an issue.

Note: On the Clariion check to see if the Failure happened at the same time as the weekly battery test, this is a scheduled activity that could effect both SPs


Just a thought.



AllBlack 397 posts since

Nov 5, 2007 2. **Re: problems after SAN failure.** Feb 23, 2009 11:42 AM

problems after SAN failure. ...

 in response to: [Lightbulb](#)

The SAN had what looks same failure last night. Haven't heard back from EMC.


This host is standalone but I was thinking of re-installing it as there were plans to add it to cluster.

The hosts in the other cluster have no dead paths as far as I can see. Although their VMs have been affected by the SAN failure. We pretty much have to reboot every VM



[Lightbulb](#) 1,391 posts since

Aug 15, 2008 3. **Re: problems after SAN failure.** Feb 23, 2009 4:11 PM


 in response to: [AllBlack](#) I suppose better safe than sorry, reinstall the host so it does not become an issue down the line.

Really hammer EMC to get an answer as what happened to the SAN things could have been much worse and you do not want that to occur again.



[AllBlack](#) 397 posts since

Nov 5, 2007 4. **Re: problems after SAN failure.** Mar 28, 2009 12:39 PM

 in response to: [Lightbulb](#)

A lot has happened since last post. Two days after that the entire SAN started to fallover and we had major outage. It was identified to a bug in Flare software that was unknown until then. Things stabilized after that and we did not use

the functionality that was buggy. A few days ago things went balls up again!!! We are getting lots of trespassing and finger


was pointed at VMware. They pretty much proofed it was SAN. It looks now that there is an issue in the hardware backend

that can cause trespassing. It has been a full-on week to say the least



[whynotq](#) 759 posts since

Mar 2, 2005 5. **Re: problems after SAN failure.** Mar 28, 2009 1:09 PM

 in response to: [AllBlack](#)

any chance of some more detail relating to the fault and environment? i work with EMC clariion constantly and have not seen a bug yet to cause this so i would be interested to here, what is the Flare code that you have currently and what is the Bug detail that EMC highlighted? did they reference any primus articles and do you have the Bug check or Panic ID?

lots of questions i know but it may help more people avoid your pain in the future 😊


i'll take a look at the sp collects if you care to post them...



[AllBlack](#) 397 posts since

Nov 5, 2007 6. **Re: problems after SAN failure.** Mar 28, 2009 1:37 PM

problems after SAN failure. ...


 in response to: [whynotq](#) I will get back to you when I have more info. They are looking into it and so is VMware. VMware has never seen anything like this. Initially EMC said it was caused by a bug in the LUN migration software so we didnt use that again.

Now after the latest round of problems, they are thinking it is all caused by a hardware fault in the backend. I don't have access to the collects right now so cannot give you more info at this time. The SAN was sending out a ASC/ASCQ 3f/0xe to ESX and apparently that points to a hardware issue. No primus reference as this looks like a first. They have engineering working on it
cheers



[whynotq](#) 759 posts since

Mar 2, 2005 7. **Re: problems after SAN failure.** Mar 28, 2009 1:50 PM

 in response to: [AilBlack](#) There are a couple of Primus articles that reference the ASC/ASCQ combination with a Sense Key of 6, they are said to be due to data changes within the LUN and do not indicate data corruption, i would expect to see these during an internal LUN migration on the Clariion, did the last migration complete?



[oroadwarrioro](#) 1 posts since

May 21, 2009 8. **Re: problems after SAN failure.** May 21, 2009 3:51 PM

 in response to: [whynotq](#)

I'm having same issue. 6 VMware ESXi servers hooked up iSCSI to Clariion AX4-5i. From time to time all Windows guests at same time get this error-

Event ID 11 - Disk - The driver detected a controller error on \Device\Harddisk1.

Event ID 15 - symmpi - The device, \Device\Scsi\symmpi1, is not ready for access yet.

Linux guests fare even worse - they either remount filesystem as read only or lock up completely.

These errors come up at exact same time from all ESX hosts in /var/log/messages-

This comes up a lot-

May 21 22:20:46 vmkernel: 1:16:47:08.634 cpu4:1305)iSCSI: bus 0 target 3 trying to establish session 0x35e402c0 to portal 0, address 10.0.0.2 port 3260 group 4

Also this-

May 21 22:19:58 vmkernel: 1:16:46:19.933 cpu7:5865)SCSI: 638: Queue for device vml.020001000060060160eaa0210086031f2471c4dd11524149442035 is being blocked to check for hung SP.

EMC replaced an SP yesterday but didn't help at all. I found this article that seemed to help a little, but issue has ocured at least once since implementing http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1008113 I used value of 32 and 16.

Did you get your issue resolved? I don't know what to do at this point.