

Austausch FC-HBA bei Sun Solaris - Teil 1

23.07.2007 Otmanix

Problemanalyse

Welcher Fehler tritt auf?

Auf dem Server „server“ (Sun Fire 6800) gibt es Probleme beim Zugriff auf das SAN. Ein FC-HBA hat auf einem Port keine Verbindung mehr zum SAN-Switch.

```
# dmesg

Thu Jul 19 12:02:49 MEST 2007
Jul  2 00:29:02 server scsi: [ID 107833 kern.warning] WARNING: /ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,37 (sd944):
Jul  2 00:29:02 server SCSI transport failed: reason 'tran_err': giving up
Jul  2 00:29:03 server scsi: [ID 107833 kern.warning] WARNING: /ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,37 (sd944):
Jul  2 00:29:03 server SCSI transport failed: reason 'tran_err': giving up
Jul  2 00:29:03 server scsi: [ID 107833 kern.warning] WARNING: /ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1f (sd938):
Jul  2 00:29:03 server SCSI transport failed: reason 'tran_err': giving up
Jul  2 00:29:03 server scsi: [ID 107833 kern.warning] WARNING: /ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1e (sd937):
Jul  2 00:29:03 server SCSI transport failed: reason 'tran_err': giving up
Jul  2 00:29:03 server scsi: [ID 107833 kern.warning] WARNING: /ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,20 (sd939):
Jul  2 00:29:03 server SCSI transport failed: reason 'tran_err': giving up
Jul  2 00:29:03 server scsi: [ID 107833 kern.warning] WARNING: /ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,38 (sd945):
Jul  2 00:29:03 server SCSI transport failed: reason 'tran_err': giving up
...
```

SCSI transport failed: reason 'tran_err': giving up sagt aus, daß ein Pfad nicht wiederhergestellt werden kann. Der betroffene Hardwarepfad ist /ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e.

```
# powermt display
Symmetrix logical device count=19
CLARiiON logical device count=0
Invista logical device count=0
=====
----- Host Bus Adapters ----- I/O Paths ----- Stats -----
### HW Path Summary Total Dead IO/Sec Q-IOs Errors
=====
2304 ssm@0/pci@1a/pci@1/lpfc@4 optimal 19 0 - 3 0
2305 ssm@0/pci@1a/pci@1/lpfc@5 failed 19 19 - 0 19
2306 ssm@0/pci@1b/pci@1/lpfc@4 optimal 19 0 - 1 0
2307 ssm@0/pci@1b/pci@1/lpfc@5 optimal 19 0 - 2 0
```

Genau das selbe meldet auch die Multipathing-Software EMC Powerpath. Dadurch, daß der Server aber über insgesamt 4 Pfade (4 Ports auf den FC-HBAs) und Multipathing verfügt, hat das keine Auswirkung auf den Betrieb.

Welche Solaris-devices sind davon betroffen?

```
# grep /ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e /etc/path_to_inst
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,9" 932 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1a" 933 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1b" 934 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1c" 935 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1d" 936 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1e" 937 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1f" 938 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,20" 939 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,21" 940 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,22" 941 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,23" 942 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,36" 943 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,37" 944 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,38" 945 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,39" 1193 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,3a" 1201 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,3b" 1202 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,3c" 1203 "sd"
"/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,3d" 1217 "sd"
# grep /ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e /etc/path_to_inst | wc -l
19
# ls -l /dev/dsk/*s0 | grep /ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d26s0 ->
../devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1a:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d27s0 ->
../devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1b:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d28s0 ->
../devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1c:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d29s0 ->
../devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1d:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d30s0 ->
../devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1e:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d31s0 ->
../devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,1f:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d32s0 ->
../devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,20:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d33s0 ->
../devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,21:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d34s0 ->
../devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,22:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d35s0 ->
../devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,23:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d54s0 ->
```

```

../././devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,36:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d55s0 ->
../././devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,37:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d56s0 ->
../././devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,38:a
lrwxrwxrwx 1 root root 59 May 3 2006 /dev/dsk/c3t30d57s0 ->
../././devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,39:a
lrwxrwxrwx 1 root other 59 May 29 12:51 /dev/dsk/c3t30d58s0 ->
../././devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,3a:a
lrwxrwxrwx 1 root other 59 May 29 12:51 /dev/dsk/c3t30d59s0 ->
../././devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,3b:a
lrwxrwxrwx 1 root other 59 May 29 12:51 /dev/dsk/c3t30d60s0 ->
../././devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,3c:a
lrwxrwxrwx 1 root other 59 May 29 17:32 /dev/dsk/c3t30d61s0 ->
../././devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,3d:a
lrwxrwxrwx 1 root root 58 May 3 2006 /dev/dsk/c3t30d9s0 ->
../././devices/ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e,9:a
# ls -l /dev/dsk/*s0 | grep /ssm@0,0/pci@1a,600000/pci@1/lpfc@5/sd@1e | wc -l

```

```

19
# powermt display dev=all | grep ssm@0/pci@1a/pci@1/lpfc@5
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d9s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d26s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d27s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d28s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d29s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d30s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d31s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d32s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d33s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d34s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d35s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d54s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d55s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d56s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d57s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d58s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d59s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d60s0 FA 1dB active dead 0 1
2305 ssm@0/pci@1a/pci@1/lpfc@5 c3t30d61s0 FA 1dB active dead 0 1
# powermt display dev=all | grep ssm@0/pci@1a/pci@1/lpfc@5 | wc -l
19

```

19 disk-devices sind von dem Ausfall betroffen. Diese sind am HBA-Port mit der Controller-Nummer 3 angeschlossen (c3) und greifen auf das Target 30 (t30) zu.

Der Controller benutzt den **lpfc-Treiber**, d.h. es handelt sich um einen Emulex FC-HBA.

Achtung:

Emulex bietet einen eigenen Treiber für Solaris an (lpfc). Zwischenzeitlich ist aber ab Solaris 10 ein Emulex-Solaris-Treiber integriert (emlxs), der in der Sun StorEdge SAN Foundation Software enthalten ist (Bestandteil des OS, Updates als OS-Patches). Aus historischen Gründen sind bei uns Original-Emulex-HBAs im Einsatz (keine Sun-branded OEM-Emulex) und auch der lpfc-Treiber!!!

Welche Hardware ist betroffen?

Wie gesagt, es handelt sich um einen Sun Midrange Server vom Typ SF6800 mit PCI Interfaceboard(s).

[Hier](http://docs.sun.com/source/819-5088-10/devicemapping.html#84417) (<http://docs.sun.com/source/819-5088-10/devicemapping.html#84417>) kann man das Devicemapping nachlesen. Daraus geht hervor, daß der ausgefallene Hardwarepfad

ssm@0,0/pci@1a,600000/pci@1 dem **IB7, Slot 3** entspricht.

```

# prtdiag
...

```

```

===== IO Cards =====

```

FRU Name	IO Type	Port ID	Bus Side	Slot	Bus Freq MHz	Max Bus Freq	Dev, Func	State	Name	Model
/NO/IB7/P0	PCI	26	B	0	33	33	1,0	ok	pci-pci8086,b152.0/pci108e,1000	pci-bridge
/NO/IB7/P0	PCI	26	B	0	33	33	0,0	ok	pci108e,1000-pci108e,1000.1	
/NO/IB7/P0	PCI	26	B	0	33	33	0,1	ok	SUNW,hme-pci108e,1001	SUNW,cheerio
/NO/IB7/P0	PCI	26	B	0	33	33	4,0	ok	SUNW,isptwo-pci1077,1020/sd (blo+	QLGC,ISP1040B
/NO/IB7/P0	PCI	26	B	1	33	33	2,0	ok	network-pci108e,abba.11	SUNW,pci-ce
/NO/IB7/P0	PCI	26	A	3	66	66	1,0	ok	pci-pci8086,b154.0/lpfc (scsi)	pci-bridge
/NO/IB7/P0	PCI	26	A	3	66	66	4,0	ok	lpfc-pci10df,f900/fp (fp)	LP9002
/NO/IB7/P0	PCI	26	A	3	66	66	5,0	ok	lpfc-pci10df,f900/fp (fp)	LP9002
/NO/IB7/P1	PCI	27	B	4	33	33	1,0	ok	pci-pci8086,b152.0/pci108e,1000	pci-bridge
/NO/IB7/P1	PCI	27	B	4	33	33	0,0	ok	pci108e,1000-pci108e,1000.1	
/NO/IB7/P1	PCI	27	B	4	33	33	0,1	ok	SUNW,hme-pci108e,1001	SUNW,cheerio
/NO/IB7/P1	PCI	27	B	4	33	33	4,0	ok	SUNW,isptwo-pci1077,1020/sd (blo+	QLGC,ISP1040B
/NO/IB7/P1	PCI	27	B	5	33	33	2,0	ok	network-pci108e,abba.11	SUNW,pci-ce
/NO/IB7/P1	PCI	27	B	6	33	33	3,0	ok	pci-pci8086,b154.0/lpfc (scsi)	pci-bridge
/NO/IB7/P1	PCI	27	B	6	33	33	4,0	ok	lpfc-pci10df,f900/fp (fp)	LP9002
/NO/IB7/P1	PCI	27	B	6	33	33	5,0	ok	lpfc-pci10df,f900/fp (fp)	LP9002
/NO/IB7/P1	PCI	27	A	7	66	66	1,0	ok	pci-pci8086,b154.0/lpfc (scsi)	pci-bridge
/NO/IB7/P1	PCI	27	A	7	66	66	4,0	ok	lpfc-pci10df,f900/fp (fp)	LP9002
/NO/IB7/P1	PCI	27	A	7	66	66	5,0	ok	lpfc-pci10df,f900/fp (fp)	LP9002

===== Active Boards for Domain =====

FRU Name	Board Type	Receptacle Status	Occupant Status	Condition Info
----------	------------	-------------------	-----------------	----------------

```
-----  
/N0/SB1  CPU_V2      connected  configured  ok          powered-on, assigned  
/N0/SB3  CPU_V2      connected  configured  ok          powered-on, assigned  
/N0/SB5  CPU_V2      connected  configured  ok          powered-on, assigned  
/N0/IB7  PCI_I/O_Boa connected  configured  ok          powered-on, assigned
```

Ein Blick in die Hardwarekonfiguration verrät, daß tatsächlich im Slot 3, 6 und 7 im IB7 Emulex LP9002-DC FC-HBAs stecken.