

## Re: lpfc: system freezing if FC connection is broken under load

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-08/0859.html>

---

**From:** Bodo Stroesser ([bstroesser\\_at\\_fujitsu-siemens.com](mailto:bstroesser_at_fujitsu-siemens.com))

**Date:** 08/03/05

Date: Wed, 03 Aug 2005 20:30:35 +0200

To: [linux-kernel@vger.kernel.org](mailto:linux-kernel@vger.kernel.org)

Bodo Stroesser wrote:

> *Hi,*  
>  
> *my dual Xeon machine freezes, if connection between*  
> *FC switch and tape drives is broken while writing to tapes.*  
>  
> *There is one SCSI target with 16 tape LUNs connected to my*  
> *FC controller via FC switch. I can reproduce the problem by*  
> *starting "dd if=/dev/zero of=/dev/st[0-7] bs=256K" on the*  
> *first 8 LUNs. Then I unplug the connection between switch and*  
> *tapes.*  
>  
> *It doesn't matter if using LP9802 or one channel of LP9402DC.*  
> *The problem happens immediately after cfg\_nodev\_tmo has*  
> *run out. If nodev\_tmo is changed, time from breaking connection*  
> *to machine freezing changes accordingly.*  
>  
> *After the problem happened, even NMIs no longer are handled.*  
> *I added nmi\_watchdog=1 to cmdline and added some simple code*  
> *to nmi handler, that writes the nmi counter directly to video ram.*  
> *In case of error, nmi no longer counts (but I have no idea, how*  
> *this can happen, maybe there is some HW bug).*

Sorry, I was wrong. The machine panics but I didn't note that, because panic\_blink() doesn't work and panic messages were not written to the fb-console.

Meanwhile I could reproduce the problem with verbose logging in lpfc-driver and with a modified panic(), that writes out full log\_buf to serial port. I've attached the logging starting from lpfc insertion. I added some comments to it, to show what I was doing.

If I understand panic info correctly, the machine crashes in st\_sleep\_done[st] because of an invalid pointer STp, but the

Linux-Kernel: Re: lpfc: system freezing if FC connection is broken under load

real reason for the crash seems to be some unexpected event in lpfc, after nodev\_tmo expired.

I hope, this helps to find the problem. If other traces are needed, please let me know.

Please CC me, I'm not on the lists.

Regards

Bodo

I inserted some comments ("#####"), but did not remove or change any console messages after reloading lpfc.

```
##### lpfc reloaded #####
<4>Emulex LightPulse Fibre Channel SCSI driver 8.0.28
<6>ACPI: PCI Interrupt 0000:04:04.0[A] -> GSI 54 (level, low) -> IRQ 54
<6>scsi6 : on PCI bus 04 device 20 irq 54
<3>lpfc 0000:04:04.0: 0:1303 Link Up Event x1 received Data: x1 xf7 x8 x0
<6>ACPI: PCI Interrupt 0000:04:05.0[A] -> GSI 55 (level, low) -> IRQ 55
<6>scsi7 : on PCI bus 04 device 28 irq 55
<3>lpfc 0000:04:05.0: 1:1303 Link Up Event x1 received Data: x1 xf7 x8 x0
<6>ACPI: PCI Interrupt 0000:09:04.0[A] -> GSI 72 (level, low) -> IRQ 72
<6>scsi8 : on PCI bus 09 device 20 irq 72
<3>lpfc 0000:09:04.0: 2:1303 Link Up Event x1 received Data: x1 xf7 x8 x0
<6>ACPI: PCI Interrupt 0000:09:05.0[A] -> GSI 73 (level, low) -> IRQ 73
<6>scsi9 : on PCI bus 09 device 28 irq 73
<3>lpfc 0000:09:05.0: 3:1303 Link Up Event x1 received Data: x1 xf7 x8 x0
<6>ACPI: PCI Interrupt 0000:0a:01.0[A] -> GSI 96 (level, low) -> IRQ 96
<6>scsi10 : on PCI bus 0a device 08 irq 96
<3>lpfc 0000:0a:01.0: 4:1303 Link Up Event x1 received Data: x1 xf7 x8 x0

##### FC connection between switch and tape devices plugged #####
<5> Vendor: IBM Model: 03590E1A Rev: E32E
<5> Type: Sequential-Access ANSI SCSI revision: 03
<4>Attached scsi tape st0 at scsi10, channel 0, id 0, lun 0
<4>st0: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295
<5>Attached scsi generic sg2 at scsi10, channel 0, id 0, lun 0, type 1
<5> Vendor: IBM Model: 03590E1A Rev: E32E
<5> Type: Sequential-Access ANSI SCSI revision: 03
<4>Attached scsi tape st1 at scsi10, channel 0, id 0, lun 1
<4>st1: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295
<5>Attached scsi generic sg3 at scsi10, channel 0, id 0, lun 1, type 1
<5> Vendor: IBM Model: 03590E1A Rev: E32E
<5> Type: Sequential-Access ANSI SCSI revision: 03
<4>Attached scsi tape st2 at scsi10, channel 0, id 0, lun 2
<4>st2: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295
<5>Attached scsi generic sg4 at scsi10, channel 0, id 0, lun 2, type 1
<5> Vendor: IBM Model: 03590E1A Rev: E32E
```

Linux–Kernel: Re: lpfc: system freezing if FC connection is broken under load

<5> Type: Sequential–Access ANSI SCSI revision: 03  
<4>Attached scsi tape st3 at scsi10, channel 0, id 0, lun 3  
<4>st3: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295  
<5>Attached scsi generic sg5 at scsi10, channel 0, id 0, lun 3, type 1  
<5> Vendor: IBM Model: 03590E1A Rev: E32E  
<5> Type: Sequential–Access ANSI SCSI revision: 03  
<4>Attached scsi tape st4 at scsi10, channel 0, id 0, lun 4  
<4>st4: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295  
<5>Attached scsi generic sg6 at scsi10, channel 0, id 0, lun 4, type 1  
<5> Vendor: IBM Model: 03590E1A Rev: E32E  
<5> Type: Sequential–Access ANSI SCSI revision: 03  
<4>Attached scsi tape st5 at scsi10, channel 0, id 0, lun 5  
<4>st5: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295  
<5>Attached scsi generic sg7 at scsi10, channel 0, id 0, lun 5, type 1  
<5> Vendor: IBM Model: 03590E1A Rev: E32E  
<5> Type: Sequential–Access ANSI SCSI revision: 03  
<4>Attached scsi tape st6 at scsi10, channel 0, id 0, lun 6  
<4>st6: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295  
<5>Attached scsi generic sg8 at scsi10, channel 0, id 0, lun 6, type 1  
<5> Vendor: IBM Model: 03590E1A Rev: E32E  
<5> Type: Sequential–Access ANSI SCSI revision: 03  
<4>Attached scsi tape st7 at scsi10, channel 0, id 0, lun 7  
<4>st7: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295  
<5>Attached scsi generic sg9 at scsi10, channel 0, id 0, lun 7, type 1  
<5> Vendor: IBM Model: 03590E1A Rev: E32E  
<5> Type: Sequential–Access ANSI SCSI revision: 03  
<4>Attached scsi tape st8 at scsi10, channel 0, id 0, lun 8  
<4>st8: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295  
<5>Attached scsi generic sg10 at scsi10, channel 0, id 0, lun 8, type 1  
<5> Vendor: IBM Model: 03590E1A Rev: E32E  
<5> Type: Sequential–Access ANSI SCSI revision: 03  
<4>Attached scsi tape st9 at scsi10, channel 0, id 0, lun 9  
<4>st9: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295  
<5>Attached scsi generic sg11 at scsi10, channel 0, id 0, lun 9, type 1  
<5> Vendor: IBM Model: 03590E1A Rev: E32E  
<5> Type: Sequential–Access ANSI SCSI revision: 03  
<4>Attached scsi tape st10 at scsi10, channel 0, id 0, lun 10  
<4>st10: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295  
<5>Attached scsi generic sg12 at scsi10, channel 0, id 0, lun 10, type 1  
<5> Vendor: IBM Model: 03590E1A Rev: E32E  
<5> Type: Sequential–Access ANSI SCSI revision: 03  
<4>Attached scsi tape st11 at scsi10, channel 0, id 0, lun 11  
<4>st11: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295  
<5>Attached scsi generic sg13 at scsi10, channel 0, id 0, lun 11, type 1  
<5> Vendor: IBM Model: 03590E1A Rev: E32E  
<5> Type: Sequential–Access ANSI SCSI revision: 03  
<4>Attached scsi tape st12 at scsi10, channel 0, id 0, lun 12  
<4>st12: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295  
<5>Attached scsi generic sg14 at scsi10, channel 0, id 0, lun 12, type 1  
<5> Vendor: IBM Model: 03590E1A Rev: E32E  
<5> Type: Sequential–Access ANSI SCSI revision: 03

Linux–Kernel: Re: lpfc: system freezing if FC connection is broken under load

```
<4>Attached scsi tape st13 at scsi10, channel 0, id 0, lun 13
<4>st13: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295
<5>Attached scsi generic sg15 at scsi10, channel 0, id 0, lun 13, type 1
<5> Vendor: IBM Model: 03590E1A Rev: E32E
<5> Type: Sequential–Access ANSI SCSI revision: 03
<4>Attached scsi tape st14 at scsi10, channel 0, id 0, lun 14
<4>st14: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295
<5>Attached scsi generic sg16 at scsi10, channel 0, id 0, lun 14, type 1
<5> Vendor: IBM Model: 03590E1A Rev: E32E
<5> Type: Sequential–Access ANSI SCSI revision: 03
<4>Attached scsi tape st15 at scsi10, channel 0, id 0, lun 15
<4>st15: try direct i/o: yes (alignment 512 B), max page reachable by HBA 4294967295
<5>Attached scsi generic sg17 at scsi10, channel 0, id 0, lun 15, type 1
```

```
##### test of 8 * "dd if=/dev/zero of=/dev/st[0–7] bs=256k &" started #####
```

```
<4>st1: Block limits 1 – 2097152 bytes.
<4>st7: Block limits 1 – 2097152 bytes.
<4>st5: Block limits 1 – 2097152 bytes.
<4>st6: Block limits 1 – 2097152 bytes.
<4>st2: Block limits 1 – 2097152 bytes.
<4>st0: Block limits 1 – 2097152 bytes.
<4>st3: Block limits 1 – 2097152 bytes.
<4>st4: Block limits 1 – 2097152 bytes.
```

```
##### lpfc_log_verbose set to 0xffff, then unplugged connection #####
```

```
<6>lpfc 0000:0a:01.0: 4:0932 FIND node did xffffd NOT FOUND Data: xff
<6>lpfc 0000:0a:01.0: 4:0112 ELS command x461 received from NPORT xffffd Data: x20
<6>lpfc 0000:0a:01.0: 4:0214 RSCN received Data: x110 x4 x20100 x0
<6>lpfc 0000:0a:01.0: 4:0247 Start Discovery Timer state x20 Data: x15 xf77f85c8 x0 x0
<6>lpfc 0000:0a:01.0: 4:0117 Xmit ELS response x2 to remote NPORT xffffd Data: x0 x4
<6>lpfc 0000:0a:01.0: 4:0128 Xmit ELS ACC response tag x0 Data: x4c xffffd x0 x0 x0
<6>lpfc 0000:0a:01.0: 4:0211 DSM in event xc on NPort x10200 in state 6 Data: x7
<6>lpfc 0000:0a:01.0: 4:0904 Add NPort x10200 to 8 list Data: x0
<6>lpfc 0000:0a:01.0: 4:0212 DSM out state 7 on NPort x10200 Data: x2010008
<6>lpfc 0000:0a:01.0: 4:0247 Start Discovery Timer state x20 Data: x15 xf77f85c8 x0 x0
<6>lpfc 0000:0a:01.0: 4:0215 RSCN processed Data: x134 x0 x1 x20
<6>lpfc 0000:0a:01.0: 4:0929 FIND node DID unmapped Data: xf5c74e80 xffffc x6 x5000503
<6>lpfc 0000:0a:01.0: 4:0236 NameServer Req Data: x171 x134 x1
<6>lpfc 0000:0a:01.0: 4:0247 Start Discovery Timer state x20 Data: x15 xf77f85c8 x0 x0
<6>lpfc 0000:0a:01.0: 4:0119 Issue GEN REQ IOCB for NPORT x220000c Data: x0 x20
<6>lpfc 0000:0a:01.0: 4:0247 Start Discovery Timer state x20 Data: x15 xf77f85c8 x0 x0
<6>lpfc 0000:0a:01.0: 4:0239 Skip x10600 NameServer Rsp Data: x0 x134 x1
<6>lpfc 0000:0a:01.0: 4:0248 Cancel Discovery Timer state x20 Data: x110 x0 x0
<6>lpfc 0000:0a:01.0: 4:0247 Start Discovery Timer state x20 Data: x15 xf77f85c8 x0 x0
<6>lpfc 0000:0a:01.0: 4:0202 Start Discovery hba state x20 Data: x134 x0 x0
```

```
##### Here nothing happens for about 30 seconds, until Nodev timeout happens #####
```

```
<3>lpfc 0000:0a:01.0: 4:0203 Nodev timeout on NPort x10200 Data: x2010008 x7 x4
<6>lpfc 0000:0a:01.0: 4:0211 DSM in event xb on NPort x10200 in state 7 Data: x2000008
<6>lpfc 0000:0a:01.0: 4:0904 Add NPort x10200 to 0 list Data: x2000000
<6>lpfc 0000:0a:01.0: 4:0212 DSM out state 255 on NPort x10200 Data: x6000000
```

Re: lpfc: system freezing if FC connection is broken under load

## Linux–Kernel: Re: lpfc: system freezing if FC connection is broken under load

```
<6>lpfc 0000:0a:01.0: 4:0900 Cleanup node for NPort x10200 Data: x2000000 x7 x4
<6>lpfc 0000:0a:01.0: 4:0904 Add NPort x10200 to 9 list Data: x2000000
<6>lpfc 0000:0a:01.0: 4:0201 Abort outstanding I/O on NPort x10200 Data: x2000000 x7 x4
<6>lpfc 0000:0a:01.0: 4:0309 Mailbox cmd x14 issue Data: x20 x700 x2
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x40011a18 x374af0c0 x0 x0 x16 x0 x32005f
x29936
<4>lpfc 0000:0a:01.0: 4:0729 FCP cmd xa failed <0/6> status: x3 result: x16 Data: x32 x5f
<6>lpfc 0000:0a:01.0: 4:0710 Iodone <0/6> cmd f2acd680, error x20008 SNS x0 x0 Data: x0 x0
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x0 x0 x0 x1 x16 x4005f x6a x20e32
<3>lpfc 0000:0a:01.0: 4:0327 Rsp ring 0 error – command completion for iotag x6a not found
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x40011a18 x2916ac48 x0 x0 x16 x0 x340059
x29936
<4>lpfc 0000:0a:01.0: 4:0729 FCP cmd xa failed <0/4> status: x3 result: x16 Data: x34 x59
<6>lpfc 0000:0a:01.0: 4:0710 Iodone <0/4> cmd e9165e00, error x20008 SNS x0 x0 Data: x0 x0
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x0 x0 x0 x1 x16 x40059 x64 x20e32
<3>lpfc 0000:0a:01.0: 4:0327 Rsp ring 0 error – command completion for iotag x64 not found
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x40011a24 x377e20c0 x0 x0 x16 x0 x37005d
x29936
<4>lpfc 0000:0a:01.0: 4:0729 FCP cmd xa failed <0/0> status: x3 result: x16 Data: x37 x5d
<6>lpfc 0000:0a:01.0: 4:0710 Iodone <0/0> cmd f2acd500, error x20008 SNS x0 x0 Data: x0 x0
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x0 x0 x0 x1 x16 x4005d x68 x20e32
<3>lpfc 0000:0a:01.0: 4:0327 Rsp ring 0 error – command completion for iotag x68 not found
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x40011a24 x377e1c48 x0 x0 x16 x0 x43005b
x29936
<4>lpfc 0000:0a:01.0: 4:0729 FCP cmd xa failed <0/3> status: x3 result: x16 Data: x43 x5b
<6>lpfc 0000:0a:01.0: 4:0710 Iodone <0/3> cmd f2acd800, error x20008 SNS x0 x0 Data: x0 x0
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x0 x0 x0 x1 x16 x4005b x66 x20e32
<3>lpfc 0000:0a:01.0: 4:0327 Rsp ring 0 error – command completion for iotag x66 not found
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x40011a18 x32362c48 x0 x0 x16 x0 x44005c
x29936
<4>lpfc 0000:0a:01.0: 4:0729 FCP cmd xa failed <0/2> status: x3 result: x16 Data: x44 x5c
<6>lpfc 0000:0a:01.0: 4:0710 Iodone <0/2> cmd f2acd980, error x20008 SNS x0 x0 Data: x0 x0
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x0 x0 x0 x1 x16 x4005c x67 x20e32
<3>lpfc 0000:0a:01.0: 4:0327 Rsp ring 0 error – command completion for iotag x67 not found
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x40011a18 x323620c0 x0 x0 x16 x0 x450058
x29936
<4>lpfc 0000:0a:01.0: 4:0729 FCP cmd xa failed <0/5> status: x3 result: x16 Data: x45 x58
<6>lpfc 0000:0a:01.0: 4:0710 Iodone <0/5> cmd f029cc80, error x20008 SNS x0 x0 Data: x0 x0
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x0 x0 x0 x1 x16 x40058 x63 x20e32
<3>lpfc 0000:0a:01.0: 4:0327 Rsp ring 0 error – command completion for iotag x63 not found
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x40011a18 x2916bc48 x0 x0 x16 x0 x48005a
x29936
<4>lpfc 0000:0a:01.0: 4:0729 FCP cmd xa failed <0/1> status: x3 result: x16 Data: x48 x5a
<6>lpfc 0000:0a:01.0: 4:0710 Iodone <0/1> cmd f029ce00, error x20008 SNS x0 x0 Data: x0 x0
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x0 x0 x0 x1 x16 x4005a x65 x20e32
<3>lpfc 0000:0a:01.0: 4:0327 Rsp ring 0 error – command completion for iotag x65 not found
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x40011a18 x374af870 x0 x0 x16 x0 x4a005e
x29936
<4>lpfc 0000:0a:01.0: 4:0729 FCP cmd xa failed <0/7> status: x3 result: x16 Data: x4a x5e
<6>lpfc 0000:0a:01.0: 4:0710 Iodone <0/7> cmd e9165c80, error x20008 SNS x0 x0 Data: x0 x0
<4>lpfc 0000:0a:01.0: 4:0326 Rsp Ring 0 error: IOCB Data: x0 x0 x0 x1 x16 x4005e x69 x20e32
```

Linux–Kernel: Re: lpfc: system freezing if FC connection is broken under load

```
<3>lpfc 0000:0a:01.0: 4:0327 Rsp ring 0 error – command completion for iotag x69 not found
<1>Unable to handle kernel paging request at virtual address 0001012d
<1> printing eip:
<4>f8b56383
<1>*pde = 00000000
<1>Oops: 0000 [#1]
<4>SMP
<4>Modules linked in: lpfc edd joydev sg st sr_mod ide_cd cdrom nvram usbserial parport_pc lp parport
thermal processor fan button ipv6 battery ac tg3 i2c_i801 i2c_core ehci_hcd uhci_hcd evdev usbcore
capability commoncap dm_mod scsi_transport_fc ext3 jbd megaraid_mbox megaraid_mm sd_mod scsi_mod
<4>CPU: 0
<4>EIP: 0060:[<f8b56383>] Not tainted VLI
<4>EFLAGS: 00010202 (2.6.12.2)
<4>EIP is at st_sleep_done+0x13/0x40 [st]
<4>eax: f2acd680 ebx: 00010101 ecx: 00000000 edx: 00020008
<4>esi: f2acd79c edi: e8fd4868 ebp: f2acd680 esp: f347ff5c
<4>ds: 007b es: 007b ss: 0068
<4>Process hotplug (pid: 19031, threadinfo=f347e000 task=eb874020)
<4>Stack: f77f8000 f8847b8c 00002002 f2acd680 f347ff78 00000000 f8847abe e9165e18
<4> e9165c98 00000003 c03a4fa0 c03ee0c0 c0124270 0000000a 00000046 00000060
<4> 00000003 bfabd888 c0124325 f347ffbc c01069ad b7f1dbb8 080bea33 c0104f1a
<4>Call Trace:
<4> [<f8847b8c>] scsi_finish_command+0x8c/0xc0 [scsi_mod]
<4> [<f8847abe>] scsi_softirq+0xbe/0xd0 [scsi_mod]
<4> [<c0124270>] __do_softirq+0x70/0xf0
<4> [<c0124325>] do_softirq+0x35/0x40
<4> [<c01069ad>] do_IRQ+0x3d/0x60
<4> [<c0104f1a>] common_interrupt+0x1a/0x20
<4>Code: 37 cf ff 58 5a e9 ee fe ff ff 8d b4 26 00 00 00 00 8d bc 27 00 00 00 00 53 8b 90 b8 00 00 00 8b 52
3c 8b 5a 38 8b 90 4c 01 00 00 <8b> 4b 2c 89 51 20 8b 90 b8 00 00 00 c7 42 38 fe ff 00 00 8b 4b
<4> <0>Kernel panic – not syncing: Fatal exception in interrupt
<4>
```

–

To unsubscribe from this list: send the line "unsubscribe linux–kernel" in the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo–info.html>

Please read the FAQ at <http://www.tux.org/lkml/>