

PROBLEM: null pointer dereference in cfq_dispatch_requests (2.6.21-rc2 and 2.6.20)

PROBLEM: null pointer dereference in cfq_dispatch_requests (2.6.21-rc2 and 2.6.20)

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-02/msg10136.html>

- *From:* Dan Williams <dan.j.williams@xxxxxxxxxx>
 - *Date:* Wed, 28 Feb 2007 11:02:35 -0700
-

I can reliably reproduce a null pointer dereference on 2.6.20 and 2.6.21-rc2. I will keep digging to find the kernel version where this last worked, but wanted to see if there were any immediate experiments I should try.

The failure is caused by running tiobench on a MD raid6 array with 6 out of 8 disks available. The commands I issued to reproduce this are:

```
mdadm -A /dev/md0 /dev/sd[bcddefg]
mount /dev/md0 /mnt/raid
tiobench --numruns 5 --size 2048 --dir /mnt/raid
```

The filesystem is ext3. The controller is an LSI 1068. Here are the two BUG messages first 2.6.21-rc2 followed by 2.6.20. I will reply to this message with the config.

Fedora Core release 5 (Bordeaux)
Kernel 2.6.21-rc2 on an i686

```
[ 431.709022] BUG: unable to handle kernel NULL pointer dereference at virtual address 0000005c
[ 431.717993] printing eip:
[ 431.720825] c01e1e00
[ 431.723112] *pde = 32e70001
[ 431.726065] Oops: 0000 [#1]
[ 431.728997] SMP
[ 431.730922] Modules linked in: raid456 xor nfsd exportfs lockd nfs_acl sunrpc autofs4 hidp l2cap bluetooth
iptables_raw xt_policy xt_multiport ipt_ULOG ipt_TTL ipt_ttl ipt_TOS ipt_tos ipt_SAME ipt_REJECT
ipt_REDIRECT ipt_recent ipt_owner ipt_NETMAP ipt_MASQUERADE ipt_LOG ipt_iprange ipt_ECN
ipt_ecn ipt_CLUSTERIP ipt_ah ipt_addrtype xt_tcpmss xt_pkttype xt_physdev xt_NFQUEUE xt_MARK
xt_mark xt_mac xt_limit xt_length xt_helper xt_dccp xt_contrack xt_CONNMARK xt_connmark
xt_CLASSIFY xt_tcpudp xt_state iptable_nat nf_nat nf_contrack_ipv4 nf_contrack iptable_mangle
nfnetlink iptable_filter ip_tables x_tables video sbs i2c_ec dock button battery asus_acpi ac radeon drm ipv6
lp parport_pc parport floppy uhci_hcd ehci_hcd e1000 i2c_i801 sg mptsas mptscsih mptbase i2c_core
scsi_transport_sas pcspkr dm_snapshot dm_zero dm_mirror dm_mod ata_piix ata_generic libata sd_mod
scsi_mod ext3 jbd
[ 431.812682] CPU: 0
[ 431.812682] EIP: 0060:[<c01e1e00>] Not tainted VLI
[ 431.812683] EFLAGS: 00010002 (2.6.21-rc2 #4)
```

PROBLEM: null pointer dereference in cfq_dispatch_requests (2.6.21-rc2 and 2.6.20)

PROBLEM: null pointer dereference in cfq_dispatch_requests (2.6.21-rc2 and 2.6.20)

```
[ 431.825386] EIP is at cfq_dispatch_insert+0xb/0x53
[ 431.830413] eax: f6c96ec0 ebx: 00000000 ecx: c0410568 edx: 00000000
[ 431.837608] esi: f7e956a4 edi: 00000000 ebp: f6c96cc0 esp: c0491e54
[ 431.844760] ds: 007b es: 007b fs: 00d8 gs: 0000 ss: 0068
[ 431.850847] Process swapper (pid: 0, ti=c0491000 task=c03ff4c0 task.ti=c0447000)
[ 431.858360] Stack: f76ae3bc f6c96cc0 00000000 f6c96cc0 c01e1fc9 00000000 000000e7 00000000
[ 431.867165] c03ffa10 c0143123 00000000 00000000 00000004 c03ff4c0 00000000 f7e957ac
[ 431.875998] f7e956a4 f7e956a4 f7d39000 f7e956a4 c01d8767 00000001 00000046 00000000
[ 431.884656] Call Trace:
[ 431.887396] [<c01e1fc9>] cfq_dispatch_requests+0x138/0x3f0
[ 431.893274] [<c0143123>] __lock_acquire+0xb64/0xbf4
[ 431.898513] [<c01d8767>] elv_next_request+0x1a1/0x1b1
[ 431.903923] [<f8a26621>] scsi_request_fn+0x59/0x336 [scsi_mod]
[ 431.910148] [<c01dbb20>] blk_run_queue+0x37/0x63
[ 431.915100] [<f8a25561>] scsi_next_command+0x25/0x2f [scsi_mod]
[ 431.921330] [<f8a2571f>] scsi_end_request+0x9e/0xa8 [scsi_mod]
[ 431.927493] [<f8a258c0>] scsi_io_completion+0x15a/0x32b [scsi_mod]
[ 431.934113] [<f882c5fb>] sd_rw_intr+0x21b/0x245 [sd_mod]
[ 431.939787] [<c031b23a>] _spin_unlock_irqrestore+0x3e/0x4d
[ 431.945640] [<f8a213f6>] scsi_finish_command+0x84/0x8b [scsi_mod]
[ 431.952051] [<c0142166>] trace_hardirqs_on+0x116/0x158
[ 431.957446] [<c012e181>] __do_softirq+0x5a/0xe9
[ 431.962329] [<c01dc291>] blk_done_softirq+0x68/0x73
[ 431.967447] [<c012e199>] __do_softirq+0x72/0xe9
[ 431.972290] [<c0107033>] do_softirq+0x6f/0xec
[ 431.976888] [<c031b0ce>] _spin_unlock_irq+0x20/0x2c
[ 431.982064] [<c0318b1b>] __sched_text_start+0x96b/0x9f3
[ 431.987574] [<c01553a1>] handle_fasteoi_irq+0x0/0xab
[ 431.992823] [<c010716d>] do_IRQ+0xbd/0xd4
[ 431.997061] [<c0105886>] common_interrupt+0x2e/0x34
[ 432.002301] [<c0103240>] mwait_idle_with_hints+0x3b/0x3f
[ 432.007931] [<c01033b9>] cpu_idle+0xb5/0xce
[ 432.012368] [<c044ca9a>] start_kernel+0x4a5/0x4ad
[ 432.017398] [<c044c1b8>] unknown_bootoption+0x0/0x202
[ 432.022829] =====
[ 432.026511] Code: 1f e9 3b c0 c7 04 24 51 6d 3a c0 e8 43 83 f4 ff e8 77 46 f2 ff ff 4f 34 83 c4 18 5b 5e 5f
5d c3 55 57 56 89 c6 53 8b 40 0c 89 d3 <8b> 7a 5c 8b 68 04 89 d0 e8 b5 fe ff ff 8b 43 14 89 da 25 01 80
[ 432.046781] EIP: [<c01e1e00>] cfq_dispatch_insert+0xb/0x53 SS:ESP 0068:c0491e54
[ 432.054403] Kernel panic – not syncing: Fatal exception in interrupt
[ 432.060912] BUG: at arch/i386/kernel/smp.c:546 smp_call_function()
[ 432.067203] [<c0118c63>] smp_call_function+0x64/0xd0
[ 432.072473] [<c023df9a>] do_unblank_screen+0x25/0x11b
[ 432.077910] [<c0118cea>] smp_send_stop+0x1b/0x40
[ 432.082848] [<c01296cb>] panic+0x54/0xfd
[ 432.087033] [<c010639c>] die+0x202/0x236
[ 432.091222] [<c031cc58>] do_page_fault+0x507/0x5e0
[ 432.096323] [<c01716e2>] kmem_cache_free+0xa1/0xb2
[ 432.101353] [<c01716e2>] kmem_cache_free+0xa1/0xb2
[ 432.106415] [<c031c751>] do_page_fault+0x0/0x5e0
[ 432.111334] [<c031b3dc>] error_code+0x7c/0x84
[ 432.115934] [<c01e1e00>] cfq_dispatch_insert+0xb/0x53
```

PROBLEM: null pointer dereference in cfq_dispatch_requests (2.6.21-rc2 and 2.6.20)

```
[ 432.121304] [<c01e1fc9>] cfq_dispatch_requests+0x138/0x3f0
[ 432.127161] [<c0143123>] __lock_acquire+0xb64/0xbf4
[ 432.132338] [<c01d8767>] elv_next_request+0x1a1/0x1b1
[ 432.137608] [<f8a26621>] scsi_request_fn+0x59/0x336 [scsi_mod]
[ 432.143762] [<c01dbb20>] blk_run_queue+0x37/0x63
[ 432.148705] [<f8a25561>] scsi_next_command+0x25/0x2f [scsi_mod]
[ 432.154884] [<f8a2571f>] scsi_end_request+0x9e/0xa8 [scsi_mod]
[ 432.160958] [<f8a258c0>] scsi_io_completion+0x15a/0x32b [scsi_mod]
[ 432.167553] [<f882c5fb>] sd_rw_intr+0x21b/0x245 [sd_mod]
[ 432.173227] [<c031b23a>] _spin_unlock_irqrestore+0x3e/0x4d
[ 432.179073] [<f8a213f6>] scsi_finish_command+0x84/0x8b [scsi_mod]
[ 432.185546] [<c0142166>] trace_hardirqs_on+0x116/0x158
[ 432.190983] [<c012e181>] __do_softirq+0x5a/0xe9
[ 432.195744] [<c01dc291>] blk_done_softirq+0x68/0x73
[ 432.200862] [<c012e199>] __do_softirq+0x72/0xe9
[ 432.205669] [<c0107033>] do_softirq+0x6f/0xec
[ 432.210294] [<c031b0ce>] _spin_unlock_irq+0x20/0x2c
[ 432.215507] [<c0318b1b>] __sched_text_start+0x96b/0x9f3
[ 432.221060] [<c01553a1>] handle_fasteoi_irq+0x0/0xab
[ 432.226256] [<c010716d>] do_IRQ+0xbd/0xd4
[ 432.230569] [<c0105886>] common_interrupt+0x2e/0x34
[ 432.235807] [<c0103240>] mwait_idle_with_hints+0x3b/0x3f
[ 432.241534] [<c01033b9>] cpu_idle+0xb5/0xce
[ 432.245948] [<c044ca9a>] start_kernel+0x4a5/0x4ad
[ 432.250972] [<c044c1b8>] unknown_bootoption+0x0/0x202
[ 432.256346] =====
```

Fedora Core release 5 (Bordeaux)
Kernel 2.6.20 on an i686

```
[ 177.299787] BUG: unable to handle kernel NULL pointer dereference at virtual address 0000005c
[ 177.308526] printing eip:
[ 177.311287] c01de510
[ 177.313521] *pde = 34d40001
[ 177.316353] Oops: 0000 [#1]
[ 177.319202] SMP
[ 177.321107] Modules linked in: raid456 xor nfsd exportfs lockd nfs_acl sunrpc autofs4 hidp l2cap bluetooth
iptables_raw xt_policy xt_multiport ipt_ULOG ipt_TTL ipt_ttl ipt_TOS ipt_tos ipt_SAME ipt_REJECT
ipt_REDIRECT ipt_recent ipt_owner ipt_NETMAP ipt_MASQUERADE ipt_LOG ipt_iprange ipt_ECN
ipt_ecn ipt_CLUSTERIP ipt_ah ipt_addrtype xt_tcpmss xt_pkttype xt_physdev xt_NFQUEUE xt_MARK
xt_mark xt_mac xt_limit xt_length xt_helper xt_dccp xt_contrack xt_CONNMARK xt_connmark
xt_CLASSIFY xt_tcpudp xt_state iptable_nat nf_nat nf_contrack_ipv4 nf_contrack iptable_mangle
nfnetlink iptable_filter ip_tables x_tables video sbs i2c_ec dock button battery asus_acpi ac radeon drm ipv6
lp parport_pc parport e1000 uhci_hcd floppy mptsas mptscsih mptbase sg ehci_hcd scsi_transport_sas
i2c_i801 i2c_core pspkr dm_snapshot dm_zero dm_mirror dm_mod ata_piix ata_generic libata sd_mod
scsi_mod ext3 jbd
[ 177.402252] CPU: 2
[ 177.402253] EIP: 0060:[<c01de510>] Not tainted VLI
[ 177.402253] EFLAGS: 00210016 (2.6.20 #5)
[ 177.414194] EIP is at cfq_dispatch_insert+0xb/0x53
[ 177.419056] eax: f7773ec0 ebx: 00000000 ecx: f7773cc0 edx: 00000000
```

PROBLEM: null pointer dereference in cfq_dispatch_requests (2.6.21-rc2 and 2.6.20)

```
[ 177.425982] esi: f70abae0 edi: f7773cc0 ebp: 00000000 esp: f34dbc9c
[ 177.432953] ds: 007b es: 007b ss: 0068
[ 177.437127] Process tiotest (pid: 5405, ti=f34db000 task=f7efc030 task.ti=f34db000)
[ 177.444763] Stack: 00000049 f77d3b9c f7773cc0 00000000 c01de6ce c014041e f8a26806 00000082
[ 177.453456] f7efc030 ffe22d6 00000000 00000000 00000000 00000004 f7efc030 f7773cc0
[ 177.462121] 00000000 00000000 00000000 f70abae0 f7cd5800 f70abae0 c01d4fcc 00000001
[ 177.470798] Call Trace:
[ 177.473503] [<c01de6ce>] cfq_dispatch_requests+0x12d/0x466
[ 177.479223] [<c014041e>] __lock_acquire+0x9e9/0xa72
[ 177.484285] [<f8a26806>] scsi_request_fn+0x286/0x336 [scsi_mod]
[ 177.490485] [<c01d4fcc>] elv_next_request+0x1a2/0x1b2
[ 177.495766] [<f8a26806>] scsi_request_fn+0x286/0x336 [scsi_mod]
[ 177.501912] [<c0315ba8>] _spin_lock_irq+0x38/0x43
[ 177.506840] [<f8a265d9>] scsi_request_fn+0x59/0x336 [scsi_mod]
[ 177.512981] [<c01d7e7d>] blk_remove_plug+0x5a/0x66
[ 177.517983] [<c01d7ea6>] __generic_unplug_device+0x1d/0x1f
[ 177.523705] [<c01d8278>] generic_unplug_device+0x15/0x21
[ 177.529272] [<f97ee054>] unplug_slaves+0x54/0x88 [raid456]
[ 177.535013] [<c01d997a>] blk_backing_dev_unplug+0x73/0x7b
[ 177.540657] [<c0315d82>] _spin_unlock_irqrestore+0x3e/0x4d
[ 177.546382] [<c0154b26>] sync_page+0x0/0x3b
[ 177.550774] [<c013f5f4>] trace_hardirqs_on+0x12e/0x158
[ 177.556108] [<c0154b26>] sync_page+0x0/0x3b
[ 177.560471] [<c018caa5>] block_sync_page+0x31/0x32
[ 177.565449] [<c0154b59>] sync_page+0x33/0x3b
[ 177.569916] [<c0313d9e>] __wait_on_bit_lock+0x2a/0x52
[ 177.575201] [<c0154b18>] __lock_page+0x58/0x5e
[ 177.579810] [<c0139612>] wake_bit_function+0x0/0x3c
[ 177.584905] [<c0155228>] do_generic_mapping_read+0x1db/0x44f
[ 177.590911] [<c01570cb>] generic_file_aio_read+0x173/0x1a4
[ 177.596617] [<c0154930>] file_read_actor+0x0/0xdb
[ 177.601525] [<c0171b47>] do_sync_read+0xc7/0x10a
[ 177.606365] [<c01395dd>] autoremove_wake_function+0x0/0x35
[ 177.612130] [<c0171a80>] do_sync_read+0x0/0x10a
[ 177.616867] [<c01723ce>] vfs_read+0xa6/0x152
[ 177.621362] [<c0172830>] sys_read+0x41/0x67
[ 177.625794] [<c0103e24>] syscall_call+0x7/0xb
[ 177.630403] =====
[ 177.634031] Code: da 11 3b c0 c7 04 24 51 9d 39 c0 e8 c9 a1 f4 ff e8 ca 6e f2 ff ff 4f 34 83 c4 18 5b 5e 5f
5d c3 55 57 56 89 c6 53 8b 40 0c 89 d3 <8b> 7a 5c 8b 68 04 89 d0 e8 b5 fe ff ff 8b 43 14 89 da 25 01 80
[ 177.654378] EIP: [<c01de510>] cfq_dispatch_insert+0xb/0x53 SS:ESP 0068:f34dbc9c
```

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx
More majordomo info at <http://vger.kernel.org/majordomo-info.html>
Please read the FAQ at <http://www.tux.org/lkml/>