

Re: [BUG] get_rtc_time() triggers NMI watchdog in hpet_rtc_interrupt()

Re: [BUG] get_rtc_time() triggers NMI watchdog in hpet_rtc_interrupt()

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-08/msg10552.html>

- *From:* "Vegard Nossum" <vegard.nossum@xxxxxxxx>
 - *Date:* Sun, 24 Aug 2008 11:14:11 +0200
-

Hi,

On Sat, Aug 23, 2008 at 11:48 AM, Mikael Pettersson <mikpe@xxxxxxxx> wrote:

Since 2.6.27-rc1 my Core2Duo has been getting sporadic oopses from hpet_rtc_interrupt, usually during shutdown or reboot, but occasionally also early in init. Today I finally managed to capture one via a serial cable:

```
INIT: version 2.86 booting
Welcome to Fedora Core
Press 'I' to enter interactive startup.
BUG: NMI Watchdog detected LOCKUP on CPU0, ip c0117092, registers:
Modules linked in: ehci_hcd uhci_hcd usbcore
```

```
Pid: 311, comm: nash-hotplug Not tainted (2.6.27-rc4 #1)
EIP: 0060:[<c0117092>] EFLAGS: 00000097 CPU: 0
EIP is at hpet_rtc_interrupt+0x2d2/0x310
EAX: 00000000 EBX: 00000002 ECX: 00000046 EDX: 00000002
ESI: 000000a6 EDI: ffff8e25 EBP: 00000008 ESP: f7bd7f28
DS: 007b ES: 007b FS: 00d8 GS: 0033 SS: 0068
Process nash-hotplug (pid: 311, ti=f7bd6000 task=f7b70460 task.ti=f7bd6000)
Stack: f7bd7f6c c0139cc0 00000000 c035ba04 00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000 00000000 f7b845a0 00000000 00000000
00000008 c01478a8 c035bf80 f7b845a0 c035bfb0 00000008 c0148f71 00000400
```

```
Call Trace:
[<c0139cc0>] hrtimer_run_pending+0x20/0x90
[<c01478a8>] handle_IRQ_event+0x28/0x50
[<c0148f71>] handle_edge_irq+0xa1/0x120
[<c010615b>] do_IRQ+0x3b/0x70
[<c0113225>] smp_apic_timer_interrupt+0x55/0x80
[<c0103c4f>] common_interrupt+0x23/0x28
[<c02c0000>] unix_release_sock+0xc0/0x220
=====
```

```
Code: 89 44 24 18 0f b6 c2 e8 5d 74 0c 00 8b 0d d8 9c 3b c0 89 44 24 1c 8b 44 24 0c 48 89
44 24 20 e9 84 fd ff ff 90 8d 74 26 00 f3 90 <a1> 80 ba 35 c0 29 f8 83 f8 01 76 f2 e9 e1 fe ff
ff 90 8d 74 26
```

Re: [BUG] get_rtc_time() triggers NMI watchdog in hpet_rtc_interrupt()

This points to the following loop in hpet_rtc_interrupt:

```
0xc0117090 <hpet_rtc_interrupt+720>: pause
0xc0117092 <hpet_rtc_interrupt+722>: mov 0xc035ba80,%eax
0xc0117097 <hpet_rtc_interrupt+727>: sub %edi,%eax
0xc0117099 <hpet_rtc_interrupt+729>: cmp $0x1,%eax
0xc011709c <hpet_rtc_interrupt+732>: jbe 0xc0117090 <hpet_rtc_interrupt+720>
```

Note: 0xc035ba80 == &jiffies

This loop originates from asm-generic/rtc.h:get_rtc_time()

```
while (jiffies - uip_watchdog < 2*HZ/100) {
    barrier();
    cpu_relax();
}
```

Note: HZ == CONFIG_HZ == 100

The bug may not originate from the 2.6.27-rc series as I only recently enabled HPET in this machine's kernels (not due to HPET problems, it inherited its .config way back from an older machine w/o HPET).

I also just got this during shutdown:

```
Syncing hardware clock to system time BUG: NMI Watchdog detected
LOCKUP on CPU0, ip c011d922, registers:
Pid: 4181, comm: hwclock Not tainted (2.6.27-rc3-00464-g1fca254-dirty #42)
EIP: 0060:[<c011d922>] EFLAGS: 00200097 CPU: 0
EIP is at hpet_rtc_interrupt+0x282/0x2e0
EAX: 00000000 EBX: 00200096 ECX: f3990000 EDX: 00010000
ESI: 000000a6 EDI: 0004f806 EBP: f3991edc ESP: f3991e98
DS: 007b ES: 007b FS: 00d8 GS: 0033 SS: 0068
Process hwclock (pid: 4181, ti=f3990000 task=f359d340 task.ti=f3990000)
Stack: f359d340 c08621c0 00000000 f359d340 00001d12 00000000 00000000 00000000
00000000 00000000 00000000 00000000 00000000 00000000 f6b4f788 00000000
00000008 f3991ef4 c017ac08 00000000 c0862180 f6b4f788 00000008 f3991f0c
Call Trace:
[<c017ac08>] ? handle_IRQ_event+0x28/0x70
[<c017c1cf>] ? handle_edge_irq+0xaf/0x140
[<c0107138>] ? do_IRQ+0x48/0xa0
[<c036e254>] ? trace_hardirqs_off_thunk+0xc/0x18
[<c0104a3c>] ? common_interrupt+0x28/0x30
[<c03c007b>] ? tty_get_baud_rate+0x3b/0x60
[<c036e641>] ? copy_from_user+0x1/0x80
[<c01bba1f>] ? sys_select+0x5f/0x190
[<c01ba157>] ? do_vfs_ioctl+0x57/0x2b0
[<c036e244>] ? trace_hardirqs_on_thunk+0xc/0x10
[<c015bc64>] ? trace_hardirqs_on_caller+0xd4/0x160
[<c0103f3b>] ? sysenter_do_call+0x12/0x3f
```

Re: [BUG] get_rtc_time() triggers NMI watchdog in hpet_rtc_interrupt()

Re: [BUG] get_rtc_time() triggers NMI watchdog in hpet_rtc_interrupt()

```
=====
Code: 65 10 25 00 89 45 d8 0f b6 45 cc e8 59 10 25 00 89 45 dc 0f b6
45 c8 e8 4d 10 25 00 83 e8 01 89 45 e0 e9 04 fe f
f ff 66 90 f3 90 <a1> 00 1b 86 c0 29 f8 83 f8 13 76 f2 e9 01 ff ff ff
83 c0 64 89
```

-----[cut here]-----

WARNING: at /uio/arkimedes/s29/vegardno/git-working/linux-2.6/kernel/mutex.c:351
mutex_trylock+0x123/0x160()

Pid: 4181, comm: hwclock Not tainted 2.6.27-rc3-00464-g1fca254-dirty #42

```
<[c0137ddf]> warn_on_slowpath+0x4f/0x80
<[c03c856d]> ? vt_console_print+0x1dd/0x2a0
<[c01589cb]> ? trace_hardirqs_off+0xb/0x10
<[c01500fb]> ? up+0x2b/0x40
<[c066fc85]> ? _spin_lock_irqsave+0x85/0xa0
<[c01383c1]> ? release_console_sem+0x1c1/0x1f0
<[c01589cb]> ? trace_hardirqs_off+0xb/0x10
<[c0670093]> ? _spin_unlock_irqrestore+0x43/0x70
<[c01383d5]> ? release_console_sem+0x1d5/0x1f0
<[c0138763]> ? vprintk+0x163/0x3c0
<[c01051f5]> ? print_trace_address+0x45/0x50
<[c0103f3b]> ? sysenter_do_call+0x12/0x3f
<[c066da53]> mutex_trylock+0x123/0x160
<[c0670093]> ? _spin_unlock_irqrestore+0x43/0x70
<[c01694cb]> crash_kexec+0x1b/0xc0
<[c038bffe]> ? vgacon_cursor+0x16e/0x1d0
<[c03c8083]> ? set_cursor+0x53/0x70
<[c03ca20b]> ? do_unblank_screen+0xbb/0x130
<[c0105dd9]> ? die_nmi+0xb9/0x100
<[c011d922]> ? hpet_rtc_interrupt+0x282/0x2e0
<[c011d922]> ? hpet_rtc_interrupt+0x282/0x2e0
<[c0105e16]> die_nmi+0xf6/0x100
<[c011d922]> ? hpet_rtc_interrupt+0x282/0x2e0
<[c0118bc5]> nmi_watchdog_tick+0x1d5/0x1e0
<[c0106227]> do_nmi+0x97/0x2d0
<[c06704f3]> nmi_stack_correct+0x26/0x2b
<[c015007b]> ? srcu_read_lock+0x3b/0x50
<[c06700d8]> ? _spin_unlock_irq+0x18/0x60
<[c011d922]> ? hpet_rtc_interrupt+0x282/0x2e0
<[c017ac08]> handle_IRQ_event+0x28/0x70
<[c017c1cf]> handle_edge_irq+0xaf/0x140
<[c0107138]> do_IRQ+0x48/0xa0
<[c036e254]> ? trace_hardirqs_off_thunk+0xc/0x18
<[c0104a3c]> common_interrupt+0x28/0x30
<[c03c007b]> ? tty_get_baud_rate+0x3b/0x60
<[c036e641]> ? copy_from_user+0x1/0x80
<[c01bba1f]> ? sys_select+0x5f/0x190
<[c01ba157]> ? do_vfs_ioctl+0x57/0x2b0
<[c036e244]> ? trace_hardirqs_on_thunk+0xc/0x10
<[c015bc64]> ? trace_hardirqs_on_caller+0xd4/0x160
<[c0103f3b]> sysenter_do_call+0x12/0x3f
```

Re: [BUG] get_rtc_time() triggers NMI watchdog in hpet_rtc_interrupt()

Re: [BUG] get_rtc_time() triggers NMI watchdog in hpet_rtc_interrupt()

---[end trace de495b709f3b3b84]---

Kernel panic – not syncing: Aiee, killing interrupt handler!

-----[cut here]-----

WARNING: at /uio/arkimedes/s29/vegardno/git-working/linux-2.6/kernel/smp.c:332

smp_call_function_mask+0x1b1/0x1c0()

Pid: 4181, comm: hwclock Tainted: G W 2.6.27-rc3-00464-g1fca254-dirty #42

[<c0137ddf>] warn_on_slowpath+0x4f/0x80

[<c038be14>] ? vgacon_set_cursor_size+0xa4/0x120

[<c01589cb>] ? trace_hardirqs_off+0xb/0x10

[<c017dd2b>] ? __rcu_read_unlock+0x9b/0xc0

[<c01506ec>] ? __atomic_notifier_call_chain+0x3c/0x50

[<c03c856d>] ? vt_console_print+0x1dd/0x2a0

[<c015071a>] ? atomic_notifier_call_chain+0x1a/0x20

[<c037d056>] ? _raw_spin_unlock+0x46/0x80

[<c066ff87>] ? _spin_unlock+0x27/0x50

[<c03c856d>] ? vt_console_print+0x1dd/0x2a0

[<c01589cb>] ? trace_hardirqs_off+0xb/0x10

[<c01500fb>] ? up+0x2b/0x40

[<c066fc85>] ? _spin_lock_irqsave+0x85/0xa0

[<c01383c1>] ? release_console_sem+0x1c1/0x1f0

[<c01589cb>] ? trace_hardirqs_off+0xb/0x10

[<c0163371>] smp_call_function_mask+0x1b1/0x1c0

[<c01383d5>] ? release_console_sem+0x1d5/0x1f0

[<c01170e0>] ? stop_this_cpu+0x0/0x50

[<c066df98>] ? mutex_unlock+0x8/0x10

[<c01589cb>] ? trace_hardirqs_off+0xb/0x10

[<c066ded4>] ? __mutex_unlock_slowpath+0xa4/0x160

[<c066df98>] ? mutex_unlock+0x8/0x10

[<c016951d>] ? crash_kexec+0x6d/0xc0

[<c01051f5>] ? print_trace_address+0x45/0x50

[<c0103f3b>] ? sysenter_do_call+0x12/0x3f

[<c01170e0>] ? stop_this_cpu+0x0/0x50

[<c01633b0>] smp_call_function+0x30/0x60

[<c011d922>] ? hpet_rtc_interrupt+0x282/0x2e0

[<c011718e>] native_smp_send_stop+0x1e/0x70

[<c011d922>] ? hpet_rtc_interrupt+0x282/0x2e0

[<c0137ccf>] panic+0x5f/0x120

[<c011d922>] ? hpet_rtc_interrupt+0x282/0x2e0

[<c013b282>] do_exit+0x7e2/0x880

[<c03c8083>] ? set_cursor+0x53/0x70

[<c03ca20b>] ? do_unblank_screen+0xbb/0x130

[<c0105dd9>] ? die_nmi+0xb9/0x100

[<c011d922>] ? hpet_rtc_interrupt+0x282/0x2e0

[<c011d922>] ? hpet_rtc_interrupt+0x282/0x2e0

[<c0105dff>] die_nmi+0xdf/0x100

[<c011d922>] ? hpet_rtc_interrupt+0x282/0x2e0

[<c0118bc5>] nmi_watchdog_tick+0x1d5/0x1e0

[<c0106227>] do_nmi+0x97/0x2d0

[<c06704f3>] nmi_stack_correct+0x26/0x2b

[<c015007b>] ? srcu_read_lock+0x3b/0x50

[<c06700d8>] ? _spin_unlock_irq+0x18/0x60

Re: [BUG] get_rtc_time() triggers NMI watchdog in hpet_rtc_interrupt()

Re: [BUG] get_rtc_time() triggers NMI watchdog in hpet_rtc_interrupt()

```
[<c011d922>] ? hpet_rtc_interrupt+0x282/0x2e0
[<c017ac08>] handle_IRQ_event+0x28/0x70
[<c017c1cf>] handle_edge_irq+0xaf/0x140
[<c0107138>] do_IRQ+0x48/0xa0
[<c036e254>] ? trace_hardirqs_off_thunk+0xc/0x18
[<c0104a3c>] common_interrupt+0x28/0x30
[<c03c007b>] ? tty_get_baud_rate+0x3b/0x60
[<c036e641>] ? copy_from_user+0x1/0x80
[<c01bba1f>] ? sys_select+0x5f/0x190
[<c01ba157>] ? do_vfs_ioctl+0x57/0x2b0
[<c036e244>] ? trace_hardirqs_on_thunk+0xc/0x10
[<c015bc64>] ? trace_hardirqs_on_caller+0xd4/0x160
[<c0103f3b>] sysenter_do_call+0x12/0x3f
=====
---[ end trace de495b709f3b3b84 ]---
```

Vegard

--

"The animistic metaphor of the bug that maliciously sneaked in while the programmer was not looking is intellectually dishonest as it disguises that the error is the programmer's own creation."

-- E. W. Dijkstra, EWD1036

--

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxxx

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

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