

## Re: Fixes for nforce2 hard lockup, apic, io-apic, udma133 covered

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2003-12/2706.html>

---

**From:** Josh McKinney (*forming\_at\_charter.net*)

**Date:** 12/11/03

Date: Thu, 11 Dec 2003 11:23:03 -0500

To: linux-kernel@vger.kernel.org

Trying to get a grasp on the all the fixes floating around. I have been running the first "timer" patch, the two liner to mpparse.c, for about five days until I made it crash with by catting 4 drives to /dev/null. It crashed after I turned on disconnect with athcool, so that may be related, because I could crash it with disconnect off. Now I am running both of Ross's patches for 2.6 for just 10 hours, but disconnect is still enabled, so far so good.

So the consensus seems to be that Ross's timer patch and the disconnect OR delay ACK patch is the mostly \*correct\* fix? As of right now I am compiling kernels with the disconnect patch and ross's timer patch, and one with those fixes and Maciej's acpi fixes below. Should I try it with just the acpi fixes sent by Maciej or are these just general fixes?

I also tried running mptable, but the output is "hosed".

Thanks

On approximately Thu, Dec 11, 2003 at 04:15:28PM +0100, Maciej W. Rozycki wrote:

> On Thu, 11 Dec 2003, Ross Dickson wrote:

>

> > ACPI: IOAPIC (id[0x02] address[0xfec00000] global\_irq\_base[0x0])

> > IOAPIC[0]: Assigned apic\_id 2

> > IOAPIC[0]: apic\_id 2, version 17, address 0xfec00000, IRQ 0-23

> > Bus #0 is ISA

> > Int: type 3, pol 0, trig 0, bus 0, irq 0, 2-0

>

> I've browsed the relevant part of the ACPI spec and the above entry is

> incorrect. It looks like INTIN0 is now the preferred line for the 8254

> timer; at least it is the default one when using ACPI tables. This is a

> bug in Linux.

>

> > ACPI: INT\_SRC\_OVR (bus[0] irq[0x0] global\_irq[0x2] polarity[0x0] trigger[0x0])

> > Int: type 0, pol 0, trig 0, bus 0, irq 0, 2-2

Linux-Kernel: Re: Fixes for nforce2 hard lockup, apic, io-apic, udma133 covered

>  
> *Now this is an explicit entry stating the 8254 timer is connected to*  
> *INTIN2. If this is not the case, the BIOS is buggy and the solution is to*  
> *fix it. I don't consider it possible to be worked around in Linux except*  
> *maybe with a command line option added manually.*  
>  
> > *ACPI: INT\_SRC\_OVR (bus[0] irq[0x9] global\_irq[0x9] polarity[0x1] trigger[0x3])*  
> > *Int: type 0, pol 1, trig 3, bus 0, irq 9, 2-9*  
>  
> *And yet another explicit entry which has an effect on configuration as*  
> *reported below.*  
>  
> > *init IO\_APIC IRQs*  
> > *IO-APIC (apicid-pin) 2-0, 2-16, 2-17, 2-18, 2-19, 2-20, 2-21, 2-22, 2-23 not connected.*  
> > *..TIMER: vector=0x31 pin1=2 pin2=-1*  
> > *..MP-BIOS bug: 8254 timer not connected to IO-APIC pin2*  
>  
> *As reported above, the BIOS explicitly reports the timer is there.*  
>  
> > *..TIMER: Is timer irq0 connected to IOAPIC Pin0? ...*  
> > *IOAPIC[0]: Set PCI routing entry (2-0 -> 0x31 -> IRQ 0 Mode:0 Active:0)*  
> > *..TIMER check 8259 ints disabled, imr1:ff, imr2:ff*  
> > *..TIMER: works OK on apic pin0 irq0*  
>  
> *And this may be correct if the default ACPI settings reflect the actual*  
> *wiring of this board (but the BIOS says otherwise).*  
>  
> > *IRQ to pin mappings:*  
> > *IRQ0 -> 0:2-> 0:0*  
> > *[...]*  
> > *IRQ9 -> 0:9-> 0:9*  
>  
> *These two entries are wrong -- the interrupts are set up as if they were*  
> *connected to multiple I/O APIC inputs. The first entry is a result of*  
> *your hack, but the second one suggests a bug somewhere.*  
>  
> > *Finally others working with kern 2.6 earlier trialled the following patch which may provide some*  
> > *more clues:*  
> > *retrieved from:*  
> >  
> > <http://www.kernel.org/pub/linux/kernel/people/bart/2.6.0-test11-bart1/broken-out/nforce2-apic.patch>  
> >  
> > *[x86] do not wrongly override mp\_ExtINT IRQ*  
>  
> *That's a workaround to the bug in Linux I've mentioned earlier. The bug*  
> *should be fixed instead. The ACPI spec doesn't support mixed*  
> *configurations, so ExtINT is irrelevant.*  
>  
> > *Perhaps someone else could get mptable to run on their machine and send you*  
> > *the result.*  
>

Linux-Kernel: Re: Fixes for nforce2 hard lockup, apic, io-apic, udma133 covered

```
> I wanted it to compare with the ACPI table and possibly to treat as a
> reference for a workaround. Since you have no valid MP-table, there's
> nothing to do.
>
> Here's a patch that fixes a few bugs I've spotted browsing through our
> ACPI code. Please try it and report the result. I don't have a system
> with ACPI available, so I cannot verify the changes at all.
>
> The same bugs are present in 2.4 and I have a corresponding patch
> available if some wants to test the changes with that version.
>
> Maciej
>
> --
> + Maciej W. Rozycki, Technical University of Gdansk, Poland +
> +-----+
> + e-mail: macro@ds2.pg.gda.pl, PGP key available +
>
> patch-mips-2.6.0-test11-20031209-acpi-irq0-1
> diff -up --recursive --new-file linux-mips-2.6.0-test11-20031209.macro/arch/i386/kernel/mpparse.c
linux-mips-2.6.0-test11-20031209/arch/i386/kernel/mpparse.c
> --- linux-mips-2.6.0-test11-20031209.macro/arch/i386/kernel/mpparse.c 2003-11-25
04:57:01.000000000 +0000
> +++ linux-mips-2.6.0-test11-20031209/arch/i386/kernel/mpparse.c 2003-12-11 09:43:26.000000000
+0000
> @@ -940,7 +940,7 @@ void __init mp_override_legacy_irq (
> * erroneously sets the trigger to level, resulting in a HUGE
> * increase of timer interrupts!
> */
> - if((bus_irq == 0) && (global_irq == 2) && (trigger == 3))
> + if((bus_irq == 0) && (trigger == 3))
> trigger = 1;
>
> intsrc.mpc_type = MP_INTSRC;
> @@ -961,7 +961,7 @@ void __init mp_override_legacy_irq (
> * Otherwise create a new entry (e.g. global_irq == 2).
> */
> for (i = 0; i < mp_irq_entries; i++) {
> - if((mp_irqs[i].mpc_dstapic == intsrc.mpc_dstapic)
> + if((mp_irqs[i].mpc_srcbus == intsrc.mpc_srcbus)
> && (mp_irqs[i].mpc_srcbusirq == intsrc.mpc_srcbusirq)) {
> mp_irqs[i] = intsrc;
> found = 1;
> @@ -1008,9 +1008,10 @@ void __init mp_config_acpi_legacy_irqs (
> */
> for (i = 0; i < 16; i++) {
>
> - if (i == 2) continue; /* Don't connect IRQ2 */
> + if (i == 2)
> + continue; /* Don't connect IRQ2 */
>
```

Linux-Kernel: Re: Fixes for nforce2 hard lockup, apic, io-apic, udma133 covered

```
> - intsrc.mpc_irqtype = i ? mp_INT : mp_ExtINT; /* 8259A to #0 */
> + intsrc.mpc_irqtype = mp_INT;
> intsrc.mpc_srcbusirq = i; /* Identity mapped */
> intsrc.mpc_dstirq = i;
>
> -
> 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/
```

--

Josh McKinney | Webmaster: <http://ioshandangie.org>

```
-----
Linux, the choice      -o) | They that can give up essential liberty
of the GNU generation  /\  | to obtain a little temporary safety deserve
                        _\v | neither liberty or safety.
                        |                                         -Benjamin Franklin
```

-

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/>