

Re: IO-APIC on nforce2 [PATCH] + [PATCH] for nmi_debug=1 + [PATCH] for idle=C1halt, 2.6.5

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-04/5771.html>

From: Ian Kumlien (pomac_at_vapor.com)

Date: 04/27/04

To: a.verweij@student.tudelft.nl

Date: Tue, 27 Apr 2004 19:35:49 +0200

On Tue, 2004-04-27 at 19:02, Arjen Verweij wrote:

> Hello,

>

> I'm sorry for the small interlude in this thread, but I just want to get

> something clear.

Imho it was a nice summation of the situation and it might be welcome for ppl that just started reading about this.

> Basically we have a problem that is all around, except for (some) Shuttle

> boards. Noone really knows what's going on, or at least if they know they

> are not vocal about it.

Yep, and asus seems to only add support for new ram manuf in dual ddr mode.

> In comes Ross Dickson. He starts poking at the problem until he comes up

> with two patches. Near the end of 2003, an NVIDIA engineer (Allen Martin)

> states that he (or maybe NVIDIA as a whole?) has been unable to reproduce

> this weird problem with hard locks, seemingly related to APIC and IO.

> He can tell us there was a bug in a reference BIOS that NVIDIA sent out

> into the world, but that it has been fixed in a follow-up. Somewhere at

> the start of December, Shuttle updates its BIOS for the AN35. Jesse Allen

> flashes the new BIOS into his board and for reasons unknown his hard lock

> problem has vanished. The importance of the update of NVIDIA's reference

> BIOS in relation to the Shuttle update of the BIOS for their product(s) is

> unknown as well.

> Meanwhile, Ross Dickson drops requests for support tickets at AMD and

> NVIDIA. Until this day, no reply yet. Unaffected by the deafening silence

> he keeps improving his patches which seem to work(tm).

Yep, and we are all great full for that =), thanks Ross.

- > *Without Ross' hard labor one can avoid the hard locks by banning APIC support from the kernel, or turn off the C1 disconnect feature in the BIOS, which is misinterpreted by one ACPI developer as running the CPU "out of spec."*

Well, it gets hot... like hell.

- > *Recently Len Brown, the ACPI Linux kernel maintainer and Intel employee – can you spot the irony? – agrees to attempt to reproduce the problem.*
- > *After having his box run with `cat /dev/hda > /dev/null` for a night straight no lockup has occurred. The brand of his motherboard is Shuttle.*
- > *Did I mention irony...?*

Heh.

- > *Although this topic is primarily about nforce2 chipsets, similar problems have been reported with SiS chipsets for AMD cpus. Other chipsets capable of having the CPU disconnect include VIA KT266(A), KT333 and KT400. For linux a tool like athcool can set the bits for the disconnect and the HLT instruction. It is unconfirmed that these chipsets suffer from the same symptoms as nforce2 chipsets.*

There are several other things that can nuke machines though.

A friend has problem with dma on a intel chipset (i keep monitoring the changelogs for fixes) but he has problems getting a > 20 says uptime. (crashes faster with dma enabled)

My firewall, a VIA Samuel 2 (microitx) dies after a few hours if you enable cpu freq. But it also seems like it changes cpu speed to often.

The common denominator with my fw and my desktop is 'to often'. Which leads me to suspect that the Hz change from 100 -> 1000 could be somewhat responsible. Could it be that we just run it to often and thus worsen the impact? And C1 disconnect shouldn't be run that often imho. Neither should cpu freq.

Perhaps some throttling would have about the same affect as Ross patches (which is what his original patches did, but not to the C1 disconnect or the HLT instruction. Could it be that some kernel code isn't well adapted to the 100 -> 1000 change?)

Anyways, that my 0.2 eur

- > *Does anyone have some input on how to tackle this problem? The only things I can come up with is mailing all the motherboard manufacturers I can think of, harass NVIDIA and/or AMD some more through proper channels (i.e. file a "bug report", but I don't expect much from this, sorry Allen) or buy Len the cheapest broken nforce2 board I can find at pricewatch.com and have it shipped to his house :)*

Heh, that would be fun if he's willing to do the work/research =).

Linux-Kernel: Re: IO-APIC on nforce2 [PATCH] + [PATCH] for nmi_debug=1 + [PATCH] for idle=C1halt, 2.6.5

PS. CC, since i'm not on this list.

--

Ian Kumlien <pomac () vapor ! com> -- <http://pomac.netswarm.net>

-

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/>

- application/pgp-signature attachment: This is a digitally signed message part