

Linux-Kernel: Re: Failing back to INSANE timesource :) Time stopped today.

Re: Failing back to INSANE timesource :) Time stopped today.

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-04/3174.html>

From: Maciej W. Rozycki (*macro_at_ds2.pg.gda.pl*)

Date: 04/15/04

Date: Thu, 15 Apr 2004 16:47:06 +0200 (CEST)
To: Niclas Gustafsson <niclas.gustafsson@codesense.com>

On Wed, 14 Apr 2004, Niclas Gustafsson wrote:

> *Watching the /proc/interrupts with 10s apart after the "stop".*

>

> *[root@s151 root]# more /proc/interrupts*

> *CPU0*

> *0: 66413955 local-APIC-edge timer*

[...]

> *LOC: 67355837*

> *ERR: 0*

> *MIS: 0*

> *[root@s151 root]# more /proc/interrupts*

> *CPU0*

> *0: 66413955 local-APIC-edge timer*

[...]

> *LOC: 67379568*

> *ERR: 0*

> *MIS: 0*

This may be because buggy SMM firmware messes with the 8259A (configured for a transparent mode -- yes that rare "local-APIC-edge" mode is tricky ;-)) insanely. You've written this is an IBM box previously -- this would be no surprise. The following patch should help -- I think it's already included in the -mm series.

```
--
+ Maciej W. Rozycki, Technical University of Gdansk, Poland +
+-----+
+ e-mail: macro@ds2.pg.gda.pl, PGP key available +
patch-2.6.5-timer_ack-2
--- linux/macro/arch/i386/kernel/io_apic.c      Wed Apr 14 03:57:24 2004
+++ linux/arch/i386/kernel/io_apic.c          Thu Apr 15 14:41:10 2004
@@ -2152,6 +2152,10 @@ static inline void check_timer(void)
 {
     int pin1, pin2;
     int vector;
+    unsigned int ver;
+

```

Re: Failing back to INSANE timesource :) Time stopped today.

Linux-Kernel: Re: Failing back to INSANE timesource :) Time stopped today.

```
+     ver = apic_read(APIC_LVR);
+     ver = GET_APIC_VERSION(ver);

/*
 * get/set the timer IRQ vector:
@@ -2165,11 +2169,15 @@ static inline void check_timer(void)
 * mode for the 8259A whenever interrupts are routed
 * through I/O APICs. Also IRQ0 has to be enabled in
 * the 8259A which implies the virtual wire has to be
- * disabled in the local APIC.
+ * disabled in the local APIC. Finally timer interrupts
+ * need to be acknowledged manually in the 8259A for
+ * do_slow_timeoffset() and for the i82489DX when using
+ * the NMI watchdog.
 */
apic_write_around(APIC_LVT0, APIC_LVT_MASKED | APIC_DM_EXTINT);
init_8259A(1);
- timer_ack = 1;
+ timer_ack = !cpu_has_tsc;
+ timer_ack |= nmi_watchdog == NMI_IO_APIC && !APIC_INTEGRATED(ver);
enable_8259A_irq(0);

pin1 = find_isa_irq_pin(0, mp_INT);
@@ -2187,7 +2195,8 @@ static inline void check_timer(void)
        disable_8259A_irq(0);
        setup_nmi();
        enable_8259A_irq(0);
-        check_nmi_watchdog();
+        if (check_nmi_watchdog() < 0)
+            timer_ack = !cpu_has_tsc;
    }
    return;
}
@@ -2210,7 +2219,8 @@ static inline void check_timer(void)
    add_pin_to_irq(0, 0, pin2);
    if (nmi_watchdog == NMI_IO_APIC) {
        setup_nmi();
-        check_nmi_watchdog();
+        if (check_nmi_watchdog() < 0)
+            timer_ack = !cpu_has_tsc;
    }
    return;
}
-

```

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/>