Subject: Re: [2.6.22] negative time jump

Posted by vaverin on Tue, 31 Jul 2007 04:43:09 GMT

View Forum Message <> Reply to Message

john stultz wrote:

- > On 7/29/07, Vasily Averin <vvs@sw.ru> wrote:
- >> I've investigated why my testnode freezes. When I found that node is freezed
- >> again I've started to press Sysrq keys and noticed the following negative time jump.

>>

>> Could anybody please help me to understand the reasons of this issue?

>>

- >> --- VvS comment: some pre-history: node boot
- >> Jul 27 13:58:10 ts28 Linux version 2.6.22 (vvs@vvs.work.ve) (gcc version 3.4.6
- >> 20060404 (Red Hat 3.4.6-3)) #11 SMP Fri Jul 27 12:47:45 MSD 2007
- >> Jul 27 13:58:10 ts28 Command line: ro root=LABEL=/1 console=ttyS0,115200
- >> console=tty debug silencelevel=8 crashkernel=128M@128M acpc=noirq clocksource=tsc

>

- > clocksource=tsc?
- > I suspect you're forcing the clocksource as its not selected by
- > default. Could you provide dmesg output without that option. It might
- > shed some light as to why the clocksource isn't chosen.

Default clocksource was acpi-pm. But the node have similar behavior when I've used this clocksource. (please look at the following logs)

Originally I've investigated SATA-related issue and noticed some strange with timers. When I've reproduced situation again I've pressed Alt+Sysrq+Q keys and noticed that it shows incorrect time (it shows 431968 sec from booting but according to the serial console timestamps it should be ~445954 sec). Then you can see that time was jumped back, next timestamp is 431965 sec, 3 sec back.

- > It may very well be your TSCs are not synched or are otherwise not
- > reliable for timekeeping, and thus time is not consistent between
- > cpus.

However IMHO it cannot explain time loops (~4400 sec) that I'm observing right now.

thank you, Vasily Averin

PS. You can look at the other logs and find more details in my attachments to

PS. You can look at the other logs and find more details in my attachments to http://bugzilla.kernel.org/show_bug.cgi?id=8650

Subject: broken timer on VIA K8T800 chipset? Posted by vaverin on Thu, 02 Aug 2007 07:28:02 GMT

Vasily Averin wrote:

- > john stultz wrote:
- >> On 7/29/07, Vasily Averin <vvs@sw.ru> wrote:
- >>> I've investigated why my testnode freezes. When I found that node is freezed
- >>> again I've started to press Sysrq keys and noticed the following negative time jump.

I've found the reason of timer-related issue on the node. I've noticed that number of timer interrupts is not changed (irq 0 in /proc/interrupts). Timer interrupt is not works, jiffies is not not incremented and timers cannot awake.

xtime and wall_to_monolitic are not changed, but ktime_get() returns non-constant values because of __get_realtime_clock_ts corrects nanoseconds. Correction vaule is different for various clocksourcses: for "tsc" clocksource it is more noticeable (~4400 sec on my node). Obviously for other clocksources correction value is not too big.

When correction value overflows we can observe negative time jump.

On the other hand timers are not critical for schedulers work, it uses readtsc for timestaps and therefore node is not hangs completely.

However Now I have new question: why timer can hangs on VIA K8T800 chipset? How is possible to fix or at least workaround this issue?

I've try to disable irqbalance -- probably it will make the node more stable.

thank you, Vasily Averin