

# Re: [RFC][PATCH] new timeofday core subsystem (v.A0)

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-09/2872.html>

---

**From:** George Anzinger ([george\\_at\\_mvista.com](mailto:george_at_mvista.com))

**Date:** 09/09/04

Date: Wed, 08 Sep 2004 17:08:55 -0700

To: john stultz <[johnstul@us.ibm.com](mailto:johnstul@us.ibm.com)>

john stultz wrote:

> *On Fri, 2004-09-03 at 17:02, George Anzinger wrote:*

>

>>> *Again, monotonic\_clock() and friends are NTP adjusted, so drift caused  
>>>by inaccurate calibration shouldn't be a problem the interval timer code  
>>>should need to worry about (outside of maybe adjusting its interval time  
>>>if its always arriving late/early). If possible the timesource  
>>>calibration code should be improved, but that's icing on the cake and  
>>>isn't critical.*

>>>

>>

>> *Are you providing a way to predict what clock count provide a given time offset  
>>INCLUDING ntp? If so, cool. If not we need to get this conversion right. We  
>>will go into this more on your return.*

>

>

> *Sorry, I'm not sure what you mean. Mind expanding on the idea while my  
> brain warms back up?*

The issue is this: A user wants a timer to fire at exactly time B which is several hours later than now (time A). We need to know how to measure this time with the timer resources (not the clock as you are talking about it). Currently we do something like `delta_jiffies = timespec_to_jiffies(B - A)` and set up a jiffies timer to expire in `delta_jiffies`. At this time we "assume" in `timespec_to_jiffies()` that we `_know_` how long a jiffie is in terms of wall clock nanoseconds. We also assume (possibly incorrectly) that this number is "good enough" even with ntp messing things up. I think this means that we assume that, on the average, we have the right conversion and that any drift will be a) small and b) on the average 0 (or real close to it).

To my mind this is a few too many assumptions, but this is what we do today. If you are proposing that ntp corrections take care of systemic constant drift (which by the way, linux use to do WRT HZ=1024 division errors), we will need a way of working this back into the `timespec_to_jiffies()` code or we will not be able to do timers with any real accuracy.

Linux-Kernel: Re: [RFC][PATCH] new timeofday core subsystem (v.A0)

An alternative, and one that I prefer, but, I think, runs counter to your proposal, is to derive the timer interrupts from the clock after the ntp correction. Some hardware will have a problem with this (x86) but other hardware is rather constrained to do things this way. The PPC counter interrupt generator, for example, counts the same pulses as the clock counter.

>

--

George Anzinger [george@mvista.com](mailto:george@mvista.com)

High-res-timers: <http://sourceforge.net/projects/high-res-timers/>

Preemption patch: <http://www.kernel.org/pub/linux/kernel/people/rml>

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to [majordomo@vger.kernel.org](mailto: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/>