

## Re: [patch] voluntary-preempt-2.6.9-rc1-bk4-Q7

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-09/0674.html>

---

**From:** Ingo Molnar ([mingo\\_at\\_elte.hu](mailto:mingo_at_elte.hu))

**Date:** 09/02/04

Date: Thu, 2 Sep 2004 16:43:01 +0200

To: Mark\_H\_Johnson@raytheon.com

\* Mark\_H\_Johnson@raytheon.com <Mark\_H\_Johnson@raytheon.com> wrote:

> *The test just completed. Over 100 traces (>500 usec) in 25 minutes*  
> *of test runs.*  
>  
> *To recap – this kernel has:*  
>  
> *Downloaded linux-2.6.8.1 from*  
> <http://www.kernel.org/pub/linux/kernel/v2.6/linux-2.6.8.1.tar.bz2>  
> *Downloaded patches from*  
> <http://redhat.com/~mingo/voluntary-preempt/diff-bk-040828-2.6.8.1.bz2>  
> <http://people.redhat.com/mingo/voluntary-preempt/voluntary-preempt-2.6.9-rc1-bk4-Q7>  
> *... email saved into mark-offset-tsc-mcount.patch ...*  
> *... email saved into ens001.patch ...*

thanks for the data. There are dozens of traces that show a big latency for no algorithmic reason, in completely unlocked codepaths. In these places the CPU seems to have an inexplicable inability to run simple sequential code that has no looping potential at all.

the NMI samples show that just about any kernel code can be delayed by this phenomenon – the kernel functions that have critical sections show up by their likely frequency of use. There doesn't seem to be anything common to the functions that show these delays, other than that they have a critical section and that they are running in your workload.

so the remaining theories are:

– DMA starvation. I've never seen anything on this scale but it's pretty much the only thing interacting with a CPU's ability to execute code – besides the other CPU running in the system.

i'd not be surprised if some audio cards tried tricks to do as aggressive DMA as physically possible, even violating hw specifications – for the purpose of producing skip-free audio output. Do you have another soundcard for testing by any chance? Another

Linux-Kernel: Re: [patch] voluntary-preempt-2.6.9-rc1-bk4-Q7

option would be to try latencytest driven not by the soundcard IRQ but by /dev/rtc.

- some sort of SMM handler that is triggered on I/O ops or something. But a number of functions in the traces dont do any I/O ops (port instructions like IN or OUT) so it's hard to imagine this to be the case. An externally triggered SMM is possible too, perhaps some independent timer triggers a watchdog SMM?

it is nearly impossible for these traces to be caused by the kernel. It really has to be some hardware effect, based on the data we have so far.

Ingo

-

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/>