

# Re: [RFC PATCH 0/6] Convert all tasklets to workqueues

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-06/msg11802.html>

---

- *From:* Steven Rostedt <[rostedt@xxxxxxxxxxxx](mailto:rostedt@xxxxxxxxxxxx)>
  - *Date:* Thu, 28 Jun 2007 11:54:35 -0400 (EDT)
- 

On Thu, 28 Jun 2007, Alexey Kuznetsov wrote:

the context-switch argument i'll believe if i see numbers. You'll probably need in excess of tens of thousands of irqs/sec to even be able to measure its overhead. (workqueues are driven by nice kernel threads so there's no TLB overhead, etc.)

It was authors of the patch who were supposed to give some numbers, at least one or two, just to prove the concept. :-)

The problem is that we don't have the hardware that uses tasklets in critical ways. My original patch series had a debug print in every function (tasklet\_schedule and friends). I got a few scattered prints on all my boxes but no flooding of prints. So I can't show that this will hurt, because on my boxes it does not.

You could set realtime priority by default, not a poor nice -5. If some network adapters were killed just because I run some task with nice --22, it would be just ridiculous.

This is my fault to the patch series. I completely forgot to up the prio. My next series will include a change where the tasklet work queue will run at something like prio FIFO 98 (or maybe 99?)

This is a bit embarrassing that I forgot to do this, since I'm a real-time developer ;-)

-- Steve

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to [majordomo@xxxxxxxxxxxxxxxx](mailto:majordomo@xxxxxxxxxxxxxxxx)

Re: [RFC PATCH 0/6] Convert all tasklets to workqueues

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>