

## Re: interrupt latency and driver locking

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2003-09/0105.html>

---

**From:** Luigi Rizzo ([rizzo\\_at\\_icir.org](mailto:rizzo_at_icir.org))

**Date:** 09/20/03

Date: Sat, 20 Sep 2003 14:11:58 -0700

To: John Polstra <[jdp@polstra.com](mailto:jdp@polstra.com)>

On Sat, Sep 20, 2003 at 01:00:54PM -0700, John Polstra wrote:

> *On 20-Sep-2003 Luigi Rizzo wrote:*

>

> > *the main problem, as i see it, is that when there are PHY events you*

> > *still need to do some expensive work while holding a lock that*

> > *blocks interrupts, with very bad impact on the worst-case*

> > *response of the system.*

>

> *I agree that is a problem, but I don't think it is the main problem.*

what i mean is that i don't think the frequency of these events (once per second or lower) is a big issue, but it is the fact that when they hit they delay potentially urgent activities for possibly very long times.

In fact, the 50..300us measured by Sam are not even the worst case;

I am pretty sure that when you have an actual PHY change the delays involved can be much longer, e.g. if you do an 'ifconfig' involving a media change (or even the initial 'up'), the "em" driver does a busy wait of well above 1 second while in the interrupt to wait for the autoconfig to settle (with DEVICE\_POLLING this hits you badly). Several other drivers have similar issues.

This is why i advocate a design change so that non-real-time, time-consuming activities are run under a different lock that does not block other real-time activities (for other subsystems of course).

cheers

luigi

---

freebsd-arch@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-arch>

To unsubscribe, send any mail to "freebsd-arch-unsubscribe@freebsd.org"