

Re: softirqd

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2003-11/2558.html>

From: Richard B. Johnson (root_at_chaos.analogic.com)

Date: 11/17/03

Date: Mon, 17 Nov 2003 09:42:34 -0500 (EST)

To: kernwek jalsl <edityacomm@yahoo.com>

On Mon, 17 Nov 2003, kernwek jalsl wrote:

- > *Hi All;*
- >
- > *On my embedded board I see that the preemption*
- > *patch(Love etal.) is not of much use because of the*
- > *following reaons:*
- >
- > *1) calling the scheduler after the hardirq returns is*
- > *not of much use because the actual work of putting the*
- > *task in the run queue is done in the bottom half*
- > *executed by ksoftirqd*
- > *2) ksoftirqd is not a real time thread*
- >
- > *Will the preemption patch work better if the ksoftirqd*
- > *is made a real time thread?*
- >
- > *Another related question : has anyone thought of*
- > *introducing prioirty among tasklets? Right now*
- > *ksoftirqd treats them in a FIFO manner. Adding*
- > *priority among the various tasklets and making sure*
- > *that ksoftirqd looks ahead in the queue would ensure a*
- > *real time treatment to a real time interrupt...*
- >

What is the problem that you are attempting to fix?
Stating that x is not y does not mean anything useful.
'Real-time' related to an 'interrupt' has no meaning.

An interrupt occurs as soon as hardware demands it.
The software response to that interrupt, i.e., interrupt-latency is dependent upon the ISR setup code in the kernel plus the execution time of any higher-priority interrupts that may be occurring plus the execution time of any prior code sharing the same interrupt. By the time your ISR code gets control, the time is bounded by some worse-case execution times and real-time has no meaning.

Then your ISR code must handle the immediate requirements of the hardware as fast and efficiently as possible. Things that can be deferred may be handed over to the soft-IRQ stuff. If you have things that can't be deferred, you must not do it in the soft-IRQ (used to be called bottom-half) handler(s).

Specifically, the things that can be handled in a soft-IRQ are hardware configuration changes and such where there are no deadlines. If you have deadlines, i.e., you must get