

## Re: splxxx level?

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2004-08/0103.html>

---

**From:** M. Warner Losh (*imp\_at\_bsdimp.com*)

**Date:** 08/30/04

Date: Mon, 30 Aug 2004 11:14:28 -0600 (MDT)

To: scottl@freebsd.org

In message: <20040830.102606.130865377.imp@bsdimp.com>

"M. Warner Losh" <imp@bsdimp.com> writes:

: In message: <41334C3B.4070101@freebsd.org>

: Scott Long <scottl@freebsd.org> writes:

:: Sam wrote:

::

:: > Hello -

:: >

:: > I'm almost to testing on my AoE driver for 4.x and have

:: > a question about interrupt priority levels.

:: >

:: > There are currently three entry points into the driver:

:: >

:: > a) strategy routine

:: > b) network frame reception routine

:: > c) timer rexmit routine

:: >

:: > Any of the three can diddle with the device structure

:: > and thusly I need to ensure they're not running simultaneously.

:: > For example, the network reception can cause a buf to be completed

:: > and the rexmit timer can cause a buf to be failed.

:: >

:: > So, what kind of contexts are the callout, strategy, and

:: > network soft interrupt called in? Which splxxx will give

:: > one of them exclusive access to whatever they need?

:: >

:: > Just as a reality check -- I am thinking about this correct, right?

:: >

:: > Cheers,

:: >

:: > Sam

:: >

::

:: With 4.x, only one CPU can be in the kernel at a time. You won't have

:: to worry about multiple processes trying to get into strategy at the

:: same time and whatnot. However, you can be preempted by your interrupt

:: handler or by a timeout or by a software interrupt like the netisr. I

## freebsd-arch: Re: splxxx level?

:: don't remember if your driver is for a specific piece of hardware or if  
:: it's a generic layer that sits in between the network interface and the  
:: block layer. If it's for dedicated hardware then you'll need to define  
:: a interrupt type in bus\_setup\_intr() and use that type for the spl  
:: (i.e. INTR\_TYPE\_NET translates to splnet(), INTR\_TYPE\_BIO translates to  
:: splbio(), etc).  
::  
:: The safe way to go is to protect all of your critical code sections with  
:: the appropriate spl level regardless. spls are very cheap and can be  
:: set/reset from an interrupt context so there is little penalty in using  
:: them liberally at first and then narrowing them down later. Just make  
:: sure that you don't leak an spl references, and don't hold an spl for so  
:: long that it creates priority inversions. Since the only interrupts and  
:: timeouts that you'll likely be dealing with are network related,  
:: splnet() is probably the right one to use.  
:  
: splimp() is what you want to use, not splnet(). Yes, this is  
: confusing, but it appears to be what all the other network drivers  
: use. None of them are using splnet() that I could find. splimp() is  
: also used by the mbuf routines to protect mbuf operations.  
:  
: splnet() is a list of the software interrupts that are network  
: drivers.  
:  
: splimp() is splnet() plus the hardware interrupts, so is more  
: appropriate to block things called from the driver. Especially one  
: that's described as having timeouts. If it is a network driver, you  
: might consider using the timeout functionality in the net stack as  
: opposed to the callout functions. This makes it possible to have  
: almost the entire driver w/o doing any spls (most of the network  
: drivers in 4 don't do spl at all, except for entry points that are  
: outside the scope of the network/interrupt entry points, eg  
: device\_suspend).

Ah, just saw 'AoE' in the reply.

This likely means that you are writing a network stack that's glued into the system with the strategy routine. Inside the strategy routine, splbio() is likely what you need to hold while dealing with the struct buf that's passed into that routine. I'm not sure what spl level you need to be at to call into the network code. I'm thinking it is splimp(), but I'm not 100% sure about that. Stevens will likely be a good resource for 4.x.

You may have to define a software interrupt to pass the packets to the network code to make the spls work out correctly.

Warner

---

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

Re: splxxx level?

freebsd-arch: Re: splxxx level?

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