

Re: Patch available for shared em interrupts (Re: em, bge, network problems survey.)

Re: Patch available for shared em interrupts (Re: em, bge, network problems survey.)

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2006-10/msg00351.html>

- *From:* Frode Nordahl <frode@xxxxxxxxxxxx>
 - *Date:* Tue, 10 Oct 2006 15:09:56 +0200
-

On 6. okt. 2006, at 04.34, Kris Kennaway wrote:

On Thu, Oct 05, 2006 at 04:05:52PM -0400, Kris Kennaway wrote:

On Wed, Oct 04, 2006 at 05:14:27PM -0600, Scott Long wrote:

All,

I'm seeing some patterns here with all of the network driver problem reports, but I need more information to help narrow it down further.

I ask all of you who are having problems to take a minute to fill out this survey and return it to Kris Kennaway (on cc:) and myself.

Thanks.

1. Are you experiencing network hangs and/or "timeout" messages on the console? If yes, please provide a *_brief_* description of the problem.

OK, next question, to all em users:

If your em device is using a shared interrupt, and you are NOT experiencing timeout problems when using this device, please let me know:

Based on successful testing on a machine with shared em interrupt, the following patch should work around the problem **in that case**.

Note that this patch will not help you if you are not using the em driver, or if you are seeing the problem with non-shared em interrupt (I have investigated on such outlier, which seems to be a problem with

Re: Patch available for shared em interrupts (Re: em, bge, network problems survey.)

a particular model of em hardware and not a generic problem with the driver).

Index: if_em.c

```
=====
RCS file: /home/ncvs/src/sys/dev/em/if_em.c,v
retrieving revision 1.65.2.18
diff -u -u -r1.65.2.18 if_em.c
--- if_em.c 25 Aug 2006 12:38:26 -0000 1.65.2.18
+++ if_em.c 5 Oct 2006 22:05:45 -0000
@@ -2086,7 +2086,7 @@
taskqueue_start_threads(&adapter->tq, 1, PI_NET, "%s taskq",
device_get_nameunit(adapter->dev));
if ((error = bus_setup_intr(dev, adapter->res_interrupt,
- INTR_TYPE_NET | INTR_FAST, em_intr_fast, adapter,
+ INTR_TYPE_NET | INTR_MPSAFE, em_intr_fast, adapter,
&adapter->int_handler_tag)) != 0) {
device_printf(dev, "Failed to register fast interrupt "
"handler: %d\n", error);
```

Please let Scott and I know whether or not this patch works for you (in addition to the information previously requested, if you have not already sent it). Unfortunately it is only a workaround, but it points to an underlying problem with fast interrupt handlers on a shared irq that can be studied separately.

I tested this on one of my other systems where em0 and USB shares an interrupt, and the patch helps to remove the watchdog timeout, and makes the system usable.

Without it the system will some times not come up successfully at all, and other times it will drop off the face of the earth as soon as some network I/O in combination with disk I/O is done.

--

Frode Nordahl

freebsd-stable@xxxxxxxxxxx mailing list

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

To unsubscribe, send any mail to "freebsd-stable-unsubscribe@xxxxxxxxxxx"