

Re: em network issues

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2006-10/msg00148.html>

- *From:* "Kip Macy" <kip.macy@xxxxxxxxx>
 - *Date:* Wed, 18 Oct 2006 15:23:48 -0700
-

I have a Sun T2000 that I generally run with the em driver from as of July in order to avoid watchdog timeouts. One trivial scenario that reproduces the problem with 100% consistency is running the ghc configure script (a 20kloc shell script) over NFS. As the T2000 doesn't exactly represent "typical" PC hardware it may not be the most desirable test platform. Nonetheless, let me know if you're interested.

Thanks for looking into this issue.

-Kip

On 10/18/06, Jack Vogel <jfvogel@xxxxxxxxx> wrote:

I think there may be a few different problems going on with the em driver on 6.2 that are being lumped under the general description of network hangs. In order to solve these I need a reproducible failure, either on a system here at Intel, or someone who is willing to be a remote guinea pig :)

I need detailed reports, meaning EXACT system data, if its an OEM box, what model, what addons, a pciconf list, description of the network, and anything special that is connected with the problem occurrence. OH, and if you have a 'before and after' situation, then please give driver deltas that worked, and which failed.

I know that there are systems out there that have management hardware that can interfere on the network, it grabs certain packets as being 'management' and doesnt pass them on to the OS. Specifically packets for port 623 and 664 get 'eaten' by this hardware. There is a fix for this, you tell the portmapper to not use ports below 665, in particular:

```
sysctl net.inet.ip.portrange.lowlast 665 (default is 600)
```

So, if you have IPMI or AMT hardware, you should try this change and see if it fixes hangs.

There is also a hardware eeprom issue on systems with an 82573 type NIC on SOME systems. There is a utility to fix that, if you

Re: em network issues

have a problem, and have that NIC email me and I can send that out to you.

Lastly, our Linux crew have long believed that there are lurking issues on some AMD based systems, we have problems with these because we dont have easy access to this hardware (as you can imagine :). But we now have evidence that SOMETIMES completion on transmit descriptors is not being written back, and this causes hangs. They (the linux team) have a modified transmit cleanup algorithm that does not use the DONE bit, instead it just using the head and tail pointers. If I can get a case where someone has this kind of hardware and has hangs AND is willing to test then perhaps I can try coding something similar up.

Also, remember to let everyone know if something gets fixed :)

Cheers,

Jack

freebsd-stable@xxxxxxxxxxx mailing list
<http://lists.freebsd.org/mailman/listinfo/freebsd-stable>
To unsubscribe, send any mail to "freebsd-stable-unsubscribe@xxxxxxxxxxx"

freebsd-net@xxxxxxxxxxx mailing list
<http://lists.freebsd.org/mailman/listinfo/freebsd-net>
To unsubscribe, send any mail to "freebsd-net-unsubscribe@xxxxxxxxxxx"