

Re: 6.3-RELEASE/7.0-RC1: em(4) fails on 82571EB_SERDES

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2008-01/msg00209.html>

- *From:* Patrick Oeschger <paketix@xxxxxxxxxxx>
 - *Date:* Wed, 23 Jan 2008 09:10:54 +0100
-

On Tue, Jan 22, 2008 at 03:49:05PM -0800, Jack Vogel wrote:

On Jan 22, 2008 9:40 AM, Patrick Oeschger <paketix@xxxxxxxxxxx> wrote:

maybe i found two issues which are em(4) related...
i tested the 6.3-RELEASE on a appliance which has two on-board SFP slots
chipset: i82571eb (serdes with tbi-interface)
em4@pci8:0:0: class=0x020000 card=0x10761903 chip=0x10608086
rev=0x06 hdr=0x00
em5@pci8:0:1: class=0x020000 card=0x10761903 chip=0x10608086
rev=0x06 hdr=0x00

first issue:

on the SFP ports link status is always up – even if no fiber optic cable is attached

this problem seems to be E1000_STATUS_LU related (in serdes mode the chipset

will use the LU flag to indicate 'signal loss' which it gets signaled from the SFP module

– PCIe* Gbe Controllers Open Source Software Developer's Manual
DocID: 316080-001 (December 2006)

8.3.3 Loss of Signal/Link Status Indication

14.8.3 Mac/SerDes (TBI-Mode) Link Setup

the following patch resolved this issue during my test:

```
--- sys/dev/em/if_em.c.orig 2008-01-22 16:55:06.000000000 +0100
+++ sys/dev/em/if_em.c 2008-01-22 13:33:58.000000000 +0100
@@ -2491,9 +2494,21 @@
 {
 struct ifnet *ifp = adapter->ifp;
 device_t dev = adapter->dev;
+ bool status_lu = E1000_READ_REG(&adapter->hw, E1000_STATUS) &
+ E1000_STATUS_LU;
+
+ if ((adapter->hw.device_id == E1000_DEV_ID_82571EB_SERDES) &&
+ (!E1000_READ_REG(&adapter->hw, E1000_TXCW) &
+ E1000_TXCW_ANE)) &&
+ (E1000_READ_REG(&adapter->hw, E1000_CTRL) &
```

Re: 6.3-RELEASE/7.0-RC1: em(4) fails on 82571EB_SERDES

```
E1000_CTRL_SLU)) {
+ if (E1000_READ_REG(&adapter->hw, E1000_STATUS) &
+ E1000_STATUS_LU) {
+ status_lu = 0;
+ } else {
+ status_lu = 1;
+ }
+ }

- if (E1000_READ_REG(&adapter->hw, E1000_STATUS) &
- E1000_STATUS_LU) {
+ if (status_lu) {
if (adapter->link_active == 0) {
e1000_get_speed_and_duplex(&adapter->hw,
&adapter->link_speed, &adapter->link_duplex);
```

second issue:

during branching of the fibre optic cabling i get a few messages regarding 'coalesced states'

Jan 22 10:36:00 test kernel: em4: 2 link states coalesced

this seems to be because of link state interrupts generated by the chipset during branching of the fiber optic cable the link can change quickly multiple times

the following patch resolved this issue during my test:

...maybe the driver needs the same for EM_FAST_IRQ and POLLING(?)

```
--- sys/dev/em/if_em.c.orig 2008-01-22 16:55:06.000000000 +0100
+++ sys/dev/em/if_em.c 2008-01-22 13:33:58.000000000 +0100
@@ -1460,7 +1460,9 @@
struct adapter *adapter = arg;
struct ifnet *ifp;
uint32_t reg_icer;
-
+ bool do_update = (adapter->hw.device_id ==
+ E1000_DEV_ID_82571EB_SERDES) ?
+ FALSE : TRUE;
EM_CORE_LOCK(adapter);
ifp = adapter->ifp;

@@ -1499,7 +1501,8 @@
callout_stop(&adapter->timer);
adapter->hw.mac.get_link_status = 1;
e1000_check_for_link(&adapter->hw);
- em_update_link_status(adapter);
+ if (do_update)
+ em_update_link_status(adapter);
/* Deal with TX cruft when link lost */
em_tx_purge(adapter);
callout_reset(&adapter->timer, hz,
```

- is there any way to get 82571_SERDES working in a future release of the

Re: 6.3-RELEASE/7.0-RC1: em(4) fails on 82571EB_SERDES

em(4) driver?

input on this issue is very welcome

br

-pat-

Thanks for the info Pat. SERDES isn't FIBER, and I'm not positive, but off the top of my head I'd say that link down is not supposed to be possible, isn't it?

SERDES is like backplane connections in Bladeservers or similar things.

On the other hand I really don't think that its been at all tested in FreeBSD so I appreciate your patch and feedback, I will try and look at your code and see about getting the fixes in.

I am very busy right now so please be patient with me :)

Jack

thanks for your feedback jack

well, i guess my description was a bit misleading then, this stuff is not related to bladeservers and/or backplanes...

the SERDES/TBI interface is the interface between the ethernet chipset and the small form-factor pluggable (SFP) module – SFP is like GBIC, just with a smaller footprint. [http://en.wikipedia.org/wiki/SFP_transceiver]

The fibre optic link will go down if the fibre optic cable gets disconnected from the SFP module – the SFP tells this to the ethernet chipset using the SERDES/TBI interface which connects the SFP to the chipset.

no worries if you are very busy right now

there is no pressure from my side to put the changes into the code

i just thought this could be interesting for others as well

let me know when you had a closer look at the code

-pat-

freebsd-net@xxxxxxxxxxx mailing list

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

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