

pfSense - Bug #5913

Link cycling issue with some ix NICs

02/19/2016 08:28 PM - Chris Buechler

Status:	Resolved	Start date:	02/19/2016
Priority:	Very High	Due date:	
Assignee:	Luiz Souza	% Done:	0%
Category:	Operating System	Estimated time:	0.00 hour
Target version:	2.3	Affected Architecture:	amd64
Affected Version:	All		

Description

There's an issue with the XG-1540's 10 Gb NICs where they cycle link upon applying certain settings to the interface, like rxcsu and txcsu, among others. Some of those things are things we do in rc.linkup or things called from it. That leaves a variety of circumstances where the NIC gets stuck in a link down/up/down/up cycle. For instance if it loses link at any point past bootup, that happens. Among a variety of other reports:

https://www.reddit.com/r/PFSENSE/comments/45bcuq/10_gig_woes/

It also has some general issue where any time the link comes up, it cycles on its own, "2 link states coalesced" is always logged.

After several hours into figuring out what all causes the thing to lose link, I came to the conclusion that it's way too many different things for us to avoid the root issue in the driver.

It doesn't happen on any Intel 10G cards with SFP+ that we have around. The XG-C2758 has no such problem.

The problem does not exist in the latest v3.1.14 driver from Intel.

Associated revisions

Revision 94bb4420 - 02/25/2016 10:13 AM - Luiz Souza

Makes interface_bring_down() remove all the CARP and IP aliases from interface.

Ticket #5913

History

#1 - 02/22/2016 12:17 AM - Jim Thompson

- Assignee set to Luiz Souza
- Affected Architecture amd64 added
- Affected Architecture deleted ()

It doesn't happen on any Intel 10G cards with SFP+ that we have around. The XG-C2758 has no such problem.

XG-C2758 uses a 82599.

I don't know what other "intel 10G cards with SFP+" you tried, so I can't comment.

The device on the 1540-D identifies as:

```
[2.3-BETA][root@xg-1540-1.localdomain]/root: pciconf -l | grep ix  
ix0@pci0:3:0:0: class=0x020000 card=0x15ad15d9 chip=0x15ad8086 rev=0x00 hdr=0x00  
ix1@pci0:3:0:1: class=0x020000 card=0x15ad15d9 chip=0x15ad8086 rev=0x00 hdr=0x00
```

This is known to the driver as: IXGBE_DEV_ID_X550EM_X_10G_T

```
ixgbe_type.h:#define IXGBE_DEV_ID_X550EM_X_10G_T    0x15AD
```

(this is on the machine that you apparently upgraded to the 3.1.14 driver).

The driver in -HEAD is 3.1.13-k, (FreeBSD-10.2 is apparently the same base driver from Intel).

https://github.com/pfsense/FreeBSD-src/blob/devel/sys/dev/ixgbe/if_ix.c#L51 says:

```
char ixgbe_driver_version[] = "3.1.13-k";
```

There are a number of differences specific to IXGBE_DEV_ID_X550EM_X_10G_T between the two revisions of the base driver, especially this:

```
if (hw->device_id == IXGBE_DEV_ID_X550EM_X_10G_T) {
    /* Config MDIO clock speed before the first MDIO PHY access */
    hlreg0 = IXGBE_READ_REG(hw, IXGBE_HLREG0);
    hlreg0 &= ~IXGBE_HLREG0_MDCSPD;
    IXGBE_WRITE_REG(hw, IXGBE_HLREG0, hlreg0);
}
```

found in ixgbe_reset_hw_X550em() in 3.1.14 driver, but not in the 3.1.13k driver. There are also a lot of minor (but could be significant) changes to the way link speeds are negotiated.

Assigned to Luiz.

#2 - 02/24/2016 02:14 PM - Luiz Souza

- Priority changed from Normal to Very High

#3 - 02/25/2016 11:52 AM - Luiz Souza

Updated the ix driver to Intel version 3.1.14.

#4 - 02/26/2016 01:08 AM - Chris Buechler

- Status changed from Confirmed to Resolved

all the circumstances that failed before are good on the latest 2.3 snapshot