

## Re: panic: invalid ife->ifm\_data (0xa) in mii\_phy\_setmedia

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2006-08/msg00622.html>

---

- *From:* Ganbold <ganbold@xxxxxxxxxxxxxx>
  - *Date:* Wed, 30 Aug 2006 17:26:07 +0900
- 

Gleb, Pyun,

I'm a kind of bit confused here whose patch to choose.  
Can you guys enlighten me in this regard?

thanks,

Ganbold

Gleb Smirnoff wrote:

Pyun,

On Wed, Aug 30, 2006 at 04:30:25PM +0900, Pyun YongHyeon wrote:

P>

```
=====
P> RCS file: /home/ncvs/src/sys/dev/bge/if_bge.c,v
P> retrieving revision 1.91.2.16
P> diff -u -r1.91.2.16 if_bge.c
P> --- if_bge.c 10 Aug 2006 11:02:14 -0000 1.91.2.16
P> +++ if_bge.c 30 Aug 2006 07:20:43 -0000
P> @@ -1007,9 +1007,26 @@
P> /* Set up the PCI DMA control register. */
P> if (sc->bge_pcie) {
P> /* PCI Express bus */
P> - dma_rw_ctl = BGE_PCI_READ_CMD|BGE_PCI_WRITE_CMD |
P> - (0xf << BGE_PCIDMARWCTL_RD_WAT_SHIFT) |
P> - (0x2 << BGE_PCIDMARWCTL_WR_WAT_SHIFT);
P> + uint32_t device_ctl;
P> +
P> + /* alternative from Linux driver */
P> + #define DMA_CTRL_WRITE_PCIE_H20MARK_128 0x00180000
P> + #define DMA_CTRL_WRITE_PCIE_H20MARK_256 0x00380000
P> +
P> + dma_rw_ctl = 0x76000000; /* XXX XXX XXX */;
P> + device_ctl = pci_read_config(sc->bge_dev,
P> + BGE_PCI_CONF_DEV_CTRL, 4);
```

Re: panic: invalid ife->ifm\_data (0xa) in mii\_phy\_setmedia

```
P> + if ((device_ctl & 0x00e0) && 0) {
P> + /*
P> + * This clause is exactly what the Broadcom-supplied
P> + * Linux does; but given overall register programming
P> + * by bge(4), this larger DMA-write watermark
P> + * value causes BCM5721 chips to totally wedge.
P> + */
P> + dma_rw_ctl |= BGE_PCIDMA_RWCTL_PCIE_WRITE_WATRMARK_256;
P> + } else {
P> + dma_rw_ctl |= BGE_PCIDMA_RWCTL_PCIE_WRITE_WATRMARK_128;
P> + }
P> } else if (sc->bge_pcix) {
```

My small penny into the discussion. I was working on reviewing the difference in initializing the PCI DMA control register in Linux tg3.c and in bge(4).

I was quite lost in this stuff, and so I asked for help from David Christensen (davidch@). He has explained me all the differences in this register between chips and I have prepared the attached patch.

Since I have very small collection of bge(4) cards, I avoid to commit it. May be I will commit it after 6.2-RELEASE if several people test it on their cards and all is OK. And it will live unmerged in HEAD until 6.3-RELEASE.

---

Index: if\_bge.c

```
=====
RCS file: /home/ncvs/src/sys/dev/bge/if_bge.c,v
retrieving revision 1.139
diff -u -p -r1.139 if_bge.c
--- if_bge.c 23 Aug 2006 11:32:54 -0000 1.139
+++ if_bge.c 23 Aug 2006 15:18:22 -0000
@@ -1005,36 +1005,48 @@ bge_chipinit(struct bge_softc *sc)
BGE_MEMWIN_WRITE(sc, i, 0);
/* Set up the PCI DMA control register. */
+ dma_rw_ctl = BGE_PCIDMARWCTL_READ_CMD |
BGE_PCIDMARWCTL_WRITE_CMD;
+
+ /* Bits 23, 22. */
+ if (sc->bge_asicrev == BGE_ASICREV_BCM5700 ||
+ sc->bge_asicrev == BGE_ASICREV_BCM5701 ||
+ sc->bge_asicrev == BGE_ASICREV_BCM5714)
+ dma_rw_ctl |= BGE_PCIDMARWCTL_ASRT_ALL_BE |
+ BGE_PCIDMARWCTL_USE_MRM;
+
+ /* DMA watermarks: bits 21 - 19, 18 - 16. */
if (sc->bge_flags & BGE_FLAG_PCIE) {
- /* PCI Express bus */
```

Re: panic: invalid ife->ifm\_data (0xa) in mii\_phy\_setmedia

Re: panic: invalid ife->ifm\_data (0xa) in mii\_phy\_setmedia

```
- dma_rw_ctl = BGE_PCI_READ_CMD|BGE_PCI_WRITE_CMD |
- (0xf << BGE_PCIDMARWCTL_RD_WAT_SHIFT) |
- (0x2 << BGE_PCIDMARWCTL_WR_WAT_SHIFT);
+ /*
+ * DMA read watermark not used on PCI-E.
+ * DMA write watermark set to 128 bytes.
+ */
+ dma_rw_ctl |= (3 << BGE_PCIDMARWCTL_WR_WAT_SHIFT);
} else if (sc->bge_flags & BGE_FLAG_PCIX) {
- /* PCI-X bus */
- if (BGE_IS_5714_FAMILY(sc)) {
- dma_rw_ctl = BGE_PCI_READ_CMD|BGE_PCI_WRITE_CMD;
- dma_rw_ctl &= ~BGE_PCIDMARWCTL_ONEDMA_ATONCE; /* XXX */
- /* XXX magic values, Broadcom-supplied Linux driver */
- if (sc->bge_asicrev == BGE_ASICREV_BCM5780)
- dma_rw_ctl |= (1 << 20) | (1 << 18) | BGE_PCIDMARWCTL_ONEDMA_ATONCE;
- else
- dma_rw_ctl |= (1 << 20) | (1 << 18) | (1 << 15);
-
- } else if (sc->bge_asicrev == BGE_ASICREV_BCM5704)
+ switch (sc->bge_asicrev) {
+ case BGE_ASICREV_BCM5780:
+ /* XXX: Linux driver magic values. */
+ dma_rw_ctl |= (1 << 20) | (1 << 18) | BGE_PCIDMARWCTL_ONEDMA_ATONCE;
+ break;
+ case BGE_ASICREV_BCM5714:
+ case BGE_ASICREV_BCM5714_A0:
+ /* XXX: Linux driver magic values. */
+ dma_rw_ctl |= (1 << 20) | (1 << 18) |
+ BGE_PCIDMARWCTL_ONEDMA_ATONCE_LOCAL;
+ break;
+ case BGE_ASICREV_BCM5704:
+ /*
+ * The 5704 uses a different encoding of read/write
+ * watermarks.
+ * watermarks: 384 bytes for write and 1536 for read.
+ */
- dma_rw_ctl = BGE_PCI_READ_CMD|BGE_PCI_WRITE_CMD |
- (0x7 << BGE_PCIDMARWCTL_RD_WAT_SHIFT) |
- (0x3 << BGE_PCIDMARWCTL_WR_WAT_SHIFT);
- else
- dma_rw_ctl = BGE_PCI_READ_CMD|BGE_PCI_WRITE_CMD |
- (0x3 << BGE_PCIDMARWCTL_RD_WAT_SHIFT) |
- (0x3 << BGE_PCIDMARWCTL_WR_WAT_SHIFT) |
- (0x0F);
+ dma_rw_ctl |= (7 << BGE_PCIDMARWCTL_RD_WAT_SHIFT) |
+ (3 << BGE_PCIDMARWCTL_WR_WAT_SHIFT);
+ break;
+ default:
+ /* All other chips: 384 for write and read. */
+ dma_rw_ctl |= (3 << BGE_PCIDMARWCTL_RD_WAT_SHIFT) |
```

Re: panic: invalid ife->ifm\_data (0xa) in mii\_phy\_setmedia

Re: panic: invalid ife->ifm\_data (0xa) in mii\_phy\_setmedia

```
+ (3 << BGE_PCIDMARWCTL_WR_WAT_SHIFT);
+ }
/*
 * 5703 and 5704 need ONEDMA_AT_ONCE as a workaround
 @@ -1047,18 +1059,20 @@ bge_chipinit(struct bge_softc *sc)
 tmp = CSR_READ_4(sc, BGE_PCI_CLKCTL) & 0x1f;
 if (tmp == 0x6 || tmp == 0x7)
 dma_rw_ctl |= BGE_PCIDMARWCTL_ONEDMA_ATONCE;
+
+ /* Set bit 23 to enable PCIX hw bug fix. */
+ dma_rw_ctl |= BGE_PCIDMARWCTL_ASRT_ALL_BE;
+ }
} else
- /* Conventional PCI bus */
- dma_rw_ctl = BGE_PCI_READ_CMD|BGE_PCI_WRITE_CMD |
- (0x7 << BGE_PCIDMARWCTL_RD_WAT_SHIFT) |
- (0x7 << BGE_PCIDMARWCTL_WR_WAT_SHIFT) |
- (0x0F);
-
- if (sc->bge_asicrev == BGE_ASICREV_BCM5703 ||
- sc->bge_asicrev == BGE_ASICREV_BCM5704 ||
- sc->bge_asicrev == BGE_ASICREV_BCM5705)
- dma_rw_ctl &= ~BGE_PCIDMARWCTL_MINDMA;
+ /* Conventional PCI bus: 1024 bytes for read and write. */
+ dma_rw_ctl |= (7 << BGE_PCIDMARWCTL_RD_WAT_SHIFT) |
+ (7 << BGE_PCIDMARWCTL_WR_WAT_SHIFT);
+
+ /* Set minimum DMA only for 5700 and 5701. */
+ if (sc->bge_asicrev == BGE_ASICREV_BCM5700 ||
+ sc->bge_asicrev == BGE_ASICREV_BCM5701)
+ dma_rw_ctl |= 0xf;
+
pci_write_config(sc->bge_dev, BGE_PCI_DMA_RW_CTL, dma_rw_ctl, 4);
/*
Index: if_bgereg.h
```

---

```
RCS file: /home/ncvs/src/sys/dev/bge/if_bgereg.h,v
retrieving revision 1.52
diff -u -p -r1.52 if_bgereg.h
--- if_bgereg.h 23 Aug 2006 11:32:54 -0000 1.52
+++ if_bgereg.h 23 Aug 2006 15:16:12 -0000
@@ -290,20 +290,25 @@
#define BGE_CHIPREV_5701_AX 0x00
/* PCI DMA Read/Write Control register */
+#define BGE_PCIDMARWCTL_WRITE_CMD 0x70000000
+#define BGE_PCIDMARWCTL_READ_CMD 0x06000000
+
#define BGE_PCIDMARWCTL_MINDMA 0x000000FF
#define BGE_PCIDMARWCTL_RDADDR_BNDRY 0x00000700
#define BGE_PCIDMARWCTL_WRADDR_BNDRY 0x00003800
#define BGE_PCIDMARWCTL_ONEDMA_ATONCE 0x00004000
```

Re: panic: invalid ife->ifm\_data (0xa) in mii\_phy\_setmedia

Re: panic: invalid ife->ifm\_data (0xa) in mii\_phy\_setmedia

```
+ #define BGE_PCIDMARWCTL_ONEDMA_ATONCE_GLOBAL 0x00004000
+ #define BGE_PCIDMARWCTL_ONEDMA_ATONCE_LOCAL 0x00008000
# define BGE_PCIDMARWCTL_RD_WAT 0x00070000
- # define BGE_PCIDMARWCTL_RD_WAT_SHIFT 16
+ # define BGE_PCIDMARWCTL_RD_WAT_SHIFT 16
# define BGE_PCIDMARWCTL_WR_WAT 0x00380000
- # define BGE_PCIDMARWCTL_WR_WAT_SHIFT 19
+ # define BGE_PCIDMARWCTL_WR_WAT_SHIFT 19
# define BGE_PCIDMARWCTL_USE_MRM 0x00400000
# define BGE_PCIDMARWCTL_ASRT_ALL_BE 0x00800000
# define BGE_PCIDMARWCTL_DFLT_PCI_RD_CMD 0x0F000000
- # define BGE_PCIDMA_RWCTL_PCI_RD_CMD_SHIFT 24
+ # define BGE_PCIDMA_RWCTL_PCI_RD_CMD_SHIFT 24
# define BGE_PCIDMARWCTL_DFLT_PCI_WR_CMD 0xF0000000
- # define BGE_PCIDMA_RWCTL_PCI_WR_CMD_SHIFT 28
+ # define BGE_PCIDMA_RWCTL_PCI_WR_CMD_SHIFT 28
# define BGE_PCI_READ_BNDRY_DISABLE 0x00000000
# define BGE_PCI_READ_BNDRY_16BYTES 0x00000100
@@ -2075,9 +2080,6 @@ struct bge_status_block {
# define BGE_MEDIA_COPPER 0x00000010
# define BGE_MEDIA_FIBER 0x00000020
- # define BGE_PCI_READ_CMD 0x06000000
- # define BGE_PCI_WRITE_CMD 0x70000000
-
# define BGE_TICKS_PER_SEC 1000000
/*
```

---

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

---

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