

Re: [PATCH] ATA over Ethernet driver for 2.6.9

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-12/1286.html>

From: Jan-Benedict Glaw (jbglaw_at_lug-owl.de)

Date: 12/06/04

Date: Mon, 6 Dec 2004 17:21:53 +0100

To: linux-kernel@vger.kernel.org

On Mon, 2004-12-06 10:51:46 -0500, Ed L Cashin <ecashin@coraid.com> wrote in message <87acsrqval.fsf@coraid.com>:

- > *The included patch allows the Linux kernel to use the ATA over*
- > *Ethernet (AoE) network protocol to communicate with any block device*
- > *that handles the AoE protocol. The Coraid EtherDrive (R) Storage*
- > *Blade is the first hardware using AoE.*
- >
- > *Like IP, AoE is an ethernet-level network protocol, registered with*
- > *the IEEE. Unlike IP, AoE is not routable.*

So AoE is out of scope for many uses...

However, some comments:

- > `diff -urpN linux-2.6.9/drivers/block/Kconfig linux-2.6.9-aoe/drivers/block/Kconfig`
- > `--- linux-2.6.9/drivers/block/Kconfig 2004-11-30 08:22:33.000000000 -0500`
- > `+++ linux-2.6.9-aoe/drivers/block/Kconfig 2004-12-06 10:40:00.000000000 -0500`
- > `@@ -357,5 +357,6 @@ config LBD`
- > *bigger than 2TB. Otherwise say N.*
- >
- > `source "drivers/s390/block/Kconfig"`
- > `+source "drivers/block/aoe/Kconfig"`
- >
- > `endmenu`
- > `diff -urpN linux-2.6.9/drivers/block/aoe/Kconfig linux-2.6.9-aoe/drivers/block/aoe/Kconfig`
- > `--- linux-2.6.9/drivers/block/aoe/Kconfig 1969-12-31 19:00:00.000000000 -0500`
- > `+++ linux-2.6.9-aoe/drivers/block/aoe/Kconfig 2004-12-06 10:40:00.000000000 -0500`
- > `@@ -0,0 +1,10 @@`
- > `+#`
- > `+# ATA over Ethernet configuration`
- > `+#`
- > `+config ATA_OVER_ETH`
- > `+ tristate "ATA over Ethernet support"`
- > `+ depends on NET`
- > `+ default m`
- > `+ help`

Linux-Kernel: Re: [PATCH] ATA over Ethernet driver for 2.6.9

- > + *This driver provides Support for ATA over Ethernet block*
- > + *devices like the Coraid EtherDrive (R) Storage Blade.*

Since your config only contains a single element, just put it into the block device Kconfig.

```
> diff -urpN linux-2.6.9/drivers/block/aoe/all.h linux-2.6.9-aoe/drivers/block/aoe/all.h
> --- linux-2.6.9/drivers/block/aoe/all.h 1969-12-31 19:00:00.000000000 -0500
> +++ linux-2.6.9-aoe/drivers/block/aoe/all.h 2004-12-06 10:40:00.000000000 -0500
[...]
```

It's not April 1st today :)

```
> + #define nelem(A) (sizeof(A) / sizeof(A)[0])
```

Just use ARRAY_SIZE()

```
> + #define AOE_MAJOR 152
> + #define ROOT_PATH "/dev/etherd/"
> + #define PATHLEN (strlen(ROOT_PATH) + 8)
```

Looks strangely hardcoded...

[many typedefs deleted]

Typedefs tend to make the code harder to read and understand. Why don't you just use the struct xxx notation?

```
> diff -urpN linux-2.6.9/drivers/block/aoe/aoeblk.c linux-2.6.9-aoe/drivers/block/aoe/aoeblk.c
> --- linux-2.6.9/drivers/block/aoe/aoeblk.c 1969-12-31 19:00:00.000000000 -0500
> +++ linux-2.6.9-aoe/drivers/block/aoe/aoeblk.c 2004-12-06 10:40:00.000000000 -0500
> + static struct block_device_operations aoe_bdops = {
> +   .open = aoeblk_open,
> +   .release = aoeblk_release,
> +   .ioctl = aoeblk_ioctl,
> +   .owner = THIS_MODULE,
> + };
```

Please use C99 syntax:

```
    .open = aoeblk_open,
```

etc.

```
> + void
> + aoeblk_gdalloc(void *vp)
> + {
> +   if (gd == nil) {
```

Could be shorter like "if (!gd) {"

```
> +void
> +aoeblk_exit(void)
> +{
> + unregister_blkdev(AOE_MAJOR, DEVICE_NAME);
> +}
> +
> +int
> +aoeblk_init(void)
> +{
> + int n;
> +
> + n = register_blkdev(AOE_MAJOR, DEVICE_NAME);
> + if (n < 0) {
> + printk(KERN_ERR "aoe: aoeblk_init: can't register major\n");
> + return n;
> + }
> + return 0;
> +}
```

__exit and __init are missing here.

```
> diff -urpN linux-2.6.9/drivers/block/aoe/aoechr.c linux-2.6.9-aoe/drivers/block/aoe/aoechr.c
> --- linux-2.6.9/drivers/block/aoe/aoechr.c 1969-12-31 19:00:00.000000000 -0500
> +++ linux-2.6.9-aoe/drivers/block/aoe/aoechr.c 2004-12-06 10:40:00.000000000 -0500
> +struct file_operations aoe_fops = {
> + write: aoechr_write,
> + read: aoechr_read,
> + open: aoechr_open,
> + release: aoechr_rel,
> + owner: THIS_MODULE,
> +};
```

C99

```
> +u16
> +nhget16(uchar *p)
> +{
> + u16 n;
> +
> + n = p[0];
> + n <<= 8;
> + return n != p[1];
> +}
> +
> +u32
> +nhget32(uchar *p)
> +{
> + u32 n;
> +
```

```

> + n = nhget16(p);
> + n <<= 16;
> + return n /= nhget16(p+2);
> +}
> +
> +void
> +hinput16(uchar *p, u16 n)
> +{
> + p[1] = n;
> + p[0] = n >>= 8;
> +}
> +
> +void
> +hinput32(uchar *p, u32 n)
> +{
> + hinput16(p+2, n);
> + hinput16(p, n >>= 16);
> +}
> +
> +u16
> +lhget16(uchar *p)
> +{
> + u16 n;
> +
> + n = p[1];
> + n <<= 8;
> + return n /= p[0];
> +}
> +
> +u32
> +lhget32(uchar *p)
> +{
> + u32 n;
> +
> + n = lhget16(p+2);
> + n <<= 16;
> + return n /= lhget16(p);
> +}
> +
> +u64
> +lhget64(uchar *p)
> +{
> + u64 n;
> +
> + n = lhget32(p+4);
> + n <<= 32;
> + return n /= lhget32(p);
> +}

```

There are function available for this, look at the endianness header files.

Linux-Kernel: Re: [PATCH] ATA over Ethernet driver for 2.6.9

After all, especially keeping in mind that AoE isn't routeable, my thinking is that this had better written as a (E)NBD server process running in userspace. This way, you'd use the in-kernel NBD driver (or the ENBD which isn't in the kernel) and you the the routing stuff for free :)

MfG, JBG

```
--  
Jan-Benedict Glaw      jbglaw@lug-owl.de      . +49-172-7608481      _ O _  
"Eine Freie Meinung in einem Freien Kopf | Gegen Zensur | Gegen Krieg  _ _ O  
 fuer einen Freien Staat voll Freier BÄrger" | im Internet! | im Irak!  O O O  
ret = do_actions((curr | FREE_SPEECH) & ~(NEW_COPYRIGHT_LAW | DRM | TCPA));
```

—
To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@vger.kernel.org
More majordomo info at <http://vger.kernel.org/majordomo-info.html>
Please read the FAQ at <http://www.tux.org/lkml/>

- application/pgp-signature attachment: Digital signature