

[OT] 8-port AHCI SATA Controller?

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-01/msg11378.html>

- *From:* Sander <sander@xxxxxxxxxxxxx>
 - *Date:* Tue, 31 Jan 2006 12:53:43 +0100
-

Hello,

I'm looking for an 8-port SATA controller based on the AHCI chipset, as according to http://linux.yyz.us/sata/sata-status.html#vendor_support this chipset is completely open.

I've searched the websides of the companies which according to <http://linux.yyz.us/sata/sata-status.html#ahci> base some of their products on this chipset, but I couldn't find an 8-port controller.

I've also googled, but without success, hence this somewhat offtopic message. Although I hope the response helps others in their quest.

The question: does an 8-port AHCI based SATA controller exist? And if, where can I find it? 12, 16 or 24 ports will do too. I don't need HW raid, just JBOD.

Some background info: I have a 24-disk chassis which is currently filled with 7 disks, and will soon be filled with 12 disks. The 7 disks are spread across two Promise SX8 controllers and the onboard nVidia controller. 3 On each SX8 and one onboard. The disks have a small partition dedicated to SW raid1 and a large partition dedicated to SW raid5. One disk acts as a spare.

The performance of the SX8 is very bad because the in-kernel driver handles one request per controller to safeguard against corruption.

This is what the comment in drivers/block/sx8.c states:

- ```
/*
* SX8 hardware has a single message queue for all ATA ports.
* When this driver was written, the hardware (firmware?) would
* corrupt data eventually, if more than one request was outstanding.
* As one can imagine, having 8 ports bottlenecking on a single
* command hurts performance.
*
* Based on user reports, later versions of the hardware (firmware?)
* seem to be able to survive with more than one command queued.
```

## [OT] 8-port AHCI SATA Controller?

```
*
* Therefore, we default to the safe option --- 1 command --- but
* allow the user to increase this.
*
* SX8 should be able to support up to ~60 queued commands
* (CARM_MAX_REQ),
* but problems seem to occur when you exceed ~30, even on newer
* hardware.
*/
static int max_queue = 1;
```

In my tests with 'static int max\_queue' I tried 30, 16, 4 and 2. With 30 and 16 I get instant corruption on the fs (ext3, which remounts ro quickly on error). With 4 I need to hit it a bit harder to see corruption, and with 2 the fs is rock solid.

- o The performance with anything larger than 2 seems acceptable.
- o It was also harder to hit with larger chunk sizes for the SW raid.

I've flashed the SX8 controllers with the latest BIOS/firmware, so that's not the problem. NCQ is disabled in the SX8 BIOS, as the in-kernel driver doesn't support NCQ and others have reported instant corruption with NCQ enabled.

I've contacted Promise Support about this (including the sx8.c comment and my findings). They replied:

===

Case Solution: I have not seen this problem in any os.

We populate all the drives and then put data into all separate logical drives. It has worked fine. One problem that you might be having is that you are creating one large volume with a bunch of smaller volume. This card is not made for that.

If you need a raid card use a card with a raid engine. If your trying to software raid it your going to get errors.

I have not seen what you are talking about because we use this card for basic or ext 3 ext 2 single hard drives. Try doing this that way.

===

This nonsense reply made me quite angry and disappointed with Promise. The SX8 seemed a really nice controller for a good price if you only need JBOD support. But I guess in the end you always get what you pay for.

FWIW, this is with a pure 64bit kernel 2.6.16-rc1 on a Tyan K8SE motherboard.

Thanks a lot in advance.

[OT] 8-port AHCI SATA Controller?

## [OT] 8-port AHCI SATA Controller?

With kind regards, Sander

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>

---

- *Follow-Ups:*

- ◆ **Re: [OT] 8-port AHCI SATA Controller?**

- ◇ *From:* Joshua Kugler

- ◆ **Re: [OT] 8-port AHCI SATA Controller?**

- ◇ *From:* Lennart Sorensen

- Prev by Date: **Re: CD writing in future Linux try #2 [ was: Re: CD writing in future Linux (stirring up a hornets' nest) ]**
- Next by Date: **[CVE-2005-4639] Re: [PATCH 2.6.14] dvb: dst: Fix possible buffer overflow**
- Previous by thread: **Another Assertion failure in journal\_start()**
- Next by thread: **Re: [OT] 8-port AHCI SATA Controller?**
- Index(es):
  - ◆ **Date**
  - ◆ **Thread**