

# Re: nfs-server silent data corruption

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2008-04/msg00448.html>

---

- *From:* "Arno J. Klaassen" <[arno@xxxxxxxxxxxxxxxxxxxxxx](mailto:arno@xxxxxxxxxxxxxxxxxxxxxx)>
  - *Date:* 21 Apr 2008 23:46:52 +0200
- 

re,

Jeremy Chadwick <[koitsu@xxxxxxxxxx](mailto:koitsu@xxxxxxxxxx)> writes:

On Mon, Apr 21, 2008 at 04:52:55PM +0200, Arno J. Klaassen wrote:

Kris Kennaway <[kris@xxxxxxxxxx](mailto:kris@xxxxxxxxxx)> writes:

Uh, you're getting server-side data corruption, it could definitely be because of the memory you added.

yop, though I'm still not convinced the memory is bad (the very same Kingston ECC as the 2\*1G in use for about half a year already) :

Can you download and run memtest86 on this system, with the added 2G ECC insalled? memtest86 doesn't guarantee showing signs of memory problems, but in most cases it'll start spewing errors almost immediately.

it finished in a bit less than 3 hours without a single error/warning

I feel pretty confident all memory is fine

One thing I did notice in the motherboard manual below is something called "Hammer Configuration". It appears to default to 800MHz, but there's an "Auto" choice. Does using Auto fix anything?

Nope

Re: nfs-server silent data corruption

I added it directly to the 2nd CPU (diagram on page 9 of [http://www.tyan.com/manuals/m\\_s2895\\_101.pdf](http://www.tyan.com/manuals/m_s2895_101.pdf)) and the problem seems to be the interaction between nfe0 and powerd .... :

That board is the weirdest thing I've seen in years.

; ) I agree I lifted (?) my eye-brows the first time I saw that diagram

Two separate CPUs using a single (shared) memory controller, two separate (and different!) nVidia chipsets, a SMSC I/O controller probably used for serial and parallel I/O, two separate nVidia NICs with Marvell PHYs (yet somehow you can bridge the two NICs and PHYs?), two separate PCI-e busses (each associated with a separate nVidia chipset), two separate PCI-X busses... the list continues.

some may say "it's just four wheels, an engine and a steer", she looks different compared to most others

I know you don't need opinions at this point, but what a behemoth. I can't imagine that thing running reliably.

though it does ;) (till the day I decided she deserved a -stable upgrade and 2 more gigs ...)

- if I stop powerd, problems go away

This would imply that clock frequency stepping is somehow attributing itself to the corruption. I don't see any BIOS options for controlling things related to AMD's Cool-n-Quiet or PowerNow! feature, which is usually what handles this.

you can turn it on/off; anyway, the problem \*seems\* easy to reproduce when freq drops quickly from 2600Mhz to 1000Mhz ....  
I just inspected a few corrupted copies, but out of 10-200Mbytes just 1 byte was 0 iso \t

Re: nfs-server silent data corruption

– I let run powerd but turn of txcsun and tso4 on the interface, the problem is a lot harder to produce (if ever this gives a hint to anyone)

Possibly shared interrupts are causing problems?

don't think so; I first had two Promise TX4 cards in this box iso the Marvell 8port card; since I had problems with TX4 some time ago I first suspected them. The board is still running memtest86, but from the dmesg I posted I don't see a shared irq.

MSI/MSI-X doing something odd? Have you tried disabling MSI/MSI-X and see if it makes a difference?

MSI is disabled as is PCI-e Error reporting (or something like that)

I think you mean "MAC LAN Bridge", according to the motherboard manual. I'm not even sure what that really does; somehow trunks the two NICs together to give you the equivalent of 2000mbit of traffic? I don't know.

probably; I never tried ;) I need the second NIC for a seperate subnet

Does the corruption you see go away if you install a separate NIC (e.g. an Intel NIC) in a PCI or PCI-e slot, and disable the onboard NICs (should be "MAC LAN: Disable" on both the primary and slave)?

Don't have one available right now (for a 2U server). I will test if I do not find another solution.

Thanx, Arno

---

freebsd-stable@xxxxxxxxxxxxx mailing list

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

To unsubscribe, send any mail to "freebsd-stable-unsubscribe@xxxxxxxxxxxxx"