

Re: nfs-server silent data corruption

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2008-04/msg00430.html>

- *From:* Jeremy Chadwick <koitsu@xxxxxxxxxxx>
 - *Date:* Mon, 21 Apr 2008 08:43:33 -0700
-

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

Kris Kennaway <kris@xxxxxxxxxxx> 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.

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?

I added it directly to the 2nd CPU (diagram on page 9 of 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.

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.

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

Re: nfs-server silent data corruption

– 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.

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

Can you boot the machine in verbose mode, and put the dmesg up somewhere?

Device is :

```
nfe0@pci:0:10:0: class=0x068000 card=0x289510f1 chip=0x005710de rev=0xa3 hdr=0x00
vendor = 'Nvidia Corp'
device = 'nForce4 Ultra NVidia Network Bus Enumerator'
class = bridge
cap 01[44] = powerspec 2 supports D0 D1 D2 D3 current D0
```

(this is with the default BIOS setting " LAN Bridge Enabled", disabling that setting makes pciconf say "class = network" but does not influence my problem)

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.

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)?

--

| Jeremy Chadwick jdc at parodius.com |
| Parodius Networking <http://www.parodius.com/> |
| UNIX Systems Administrator Mountain View, CA, USA |
| Making life hard for others since 1977. PGP: 4BD6C0CB |

Re: nfs-server silent data corruption

freebsd-stable@xxxxxxxxxxx mailing list

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

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