

Sunday, October 11. 2009

Observations on memory reliability

Robin Harris points to an interesting study about DRAM failures in his blog storagemojo (BTW: Robins blog is really a great read). He points to the paper "DRAM Errors in the Wild: A Large-Scale Field Study" (written by Bianca Schroeder (University of Toronto), Eduardo Pinheiro and Wolf-Dietrich Weber (both from Google)) in his article.

Some of the numbers are really terrifying: 4.15% unrecoverable errors for of of the platforms are much more then i had thought and i'm somewhat conservative in my thinking how far i trust hardware. Furthermore hard errors (as in "bit permanently flipped and put it to the trashbin") are vastly more common reasons for errors as most people think.

As a sidenote: After the discussion about DRAM prices in the M3000 i've got some flak because of memory prices at different quality and got many comments "there never failed a dimm at my home pc". But given the point that Google is said to use cheaper hardware and the amount of errors (especially the unrecoverable ones) there may be a point behind the fixation of Sun in regard of memory quality

Posted by Joerg Moellenkamp in English, Technology, The IT Business at 20:48

The study's interesting, but drawing any conclusions is difficult as it doesn't say what memory Google buys.

There is a suggestion (without any evidence given, so...) that Google buys RAM in massive bulk that's failed mfr QA, but for absolute peanuts. Google then check the RAM themselves.

Given Google's scale, that could actually work out.

But it skews the findings of that paper a bit.

Read the discussion at arstechnica on the story.
Anonymous on Oct 11 2009, 21:51

yeah, the results show that the majority of errors appear on the same machines, those have a lot of errors, but there are enough of machines with ram that doesn't have any errors.
the conclusion is still valid, lots of bad ram out there and not only marketed to the "chinese homeland"...
something mentioned elsewhere is the conclusion that the not so good chips are used for cheap ecc modules because it doesnt matter that much there.

Anonymous on Oct 12 2009, 02:06

Hmmm ... i don't know if this would really save that much money. AFAIK the chips are tested before soldering, so defunct chips wouldn't get on the PCB. And a defunct PCB is often a total loss and you could just desolder the chips.

It's just an assumption, but: When google knocks at your door and says "I want xy Petabyte RAM" you get pretty competitive prices without buying the trash bin.

Dont't know how much truth is behind this, but i don't think that they would publish such an paper when they know that their DIMMS were the ones left after after sweeping the fab
Anonymous on Oct 12 2009, 06:54

a once-per-hour to once-per-day error rate in 8% of DIMMs is quite huge imho.
Anonymous on Oct 12 2009, 10:42

Very interesting article, thanks for sharing.
Anonymous on Oct 14 2009, 11:08