

- [Home](#)
- [Contact & About](#)
- [Price Lists](#)
- [Archive](#)
- [Advertise](#)
- [Subscribe](#)



CERN's data corruption research

by Robin Harris on Wednesday, 19 September, 2007

I was surprised at how many ZDnet readers reacted with disbelief to my recent [Storage Bits](#) series on data corruption (see [How data gets lost](#), [50 ways to lose your data](#) and [How Microsoft puts your data at risk](#)), claiming it had never happened to them.

Then I thought about it

What does data corruption look like to users? Does a window pop up with big red letters blaring "DATA CORRUPTION!!!" Nope, we get these "File not found" and other notices that could be - but who knows? - related to data corruption. Something goes badly wrong and you have to reinstall an application or the OS. But really, how prevalent is data corruption?

CERN does some research

That's why I was delighted to see a new paper from CERN. Now, finally, some statistics are in, reported in a recent paper titled [Data Integrity](#) by Bernd Panzer-Steindel of the CERN IT group.

Petabytes of on-disk data analyzed

At CERN, the world's largest particle physics lab, several researchers have analyzed the creation and propagation of silent data corruption. CERN's huge collider - built beneath Switzerland and France - will generate 15 thousand terabytes of data next year.

The experiments at CERN - high energy "shots" that create many terabytes of data in a few seconds - then require months of careful statistical analysis to find traces of rare and short-lived particles. Errors in the data could invalidate the results, so CERN scientists and engineers did a systematic analysis to find silent data corruption events.

The program

The analysis looked at data corruption at 3 levels:

- **Disk errors.** They wrote a special 2 GB file to more than 3,000 nodes every 2 hours and read it back checking for errors for 5 weeks. They found 500 errors on 100 nodes.
 - **Single bit errors.** 10% of disk errors.
 - **Sector (512 bytes) sized errors.** 10% of disk errors.
 - **64 KB regions.** 80% of disk errors. This one turned out to be a bug in WD disk firmware interacting with 3Ware controller cards which CERN fixed by updating the firmware in 3,000 drives.
- **RAID errors.** They ran the verify command on 492 RAID systems each week for 4 weeks. The disks are spec'd at a Bit Error Rate of 10^{-14} read/written. The good news is that the observed BER was only about a 3rd of the spec'd rate. The bad news is that in reading/writing 2.4 petabytes of data there were some 300 errors.
- **Memory errors.** Good news: only 3 double-bit errors in 3 months on 1300 nodes. Bad news: according to the spec there shouldn't have been any. Only double bit errors can't be corrected.

All of these errors will corrupt user data. When they checked 8.7 TB of user data for corruption - 33,700 files - they found 22 corrupted files, or 1 in every 1500 files.

The bottom line

CERN found an overall byte error rate of $3 * 10^{-7}$, a rate considerably higher than numbers like 10^{-14} or 10^{-12} spec'd

for components would suggest. This isn't sinister.

It's the BER of each link in the chain from CPU to disk and back again plus the fact that for some traffic, such as transferring a byte from the network to a disk, requires 6 memory r/w operations. That really pumps up the data volume and with it the likelihood of encountering an error.

The cost of accuracy

Accuracy isn't free. The CERN paper concludes that taking measures to improve accuracy

. . . will lead to a doubling of the original required IO performance on the disk servers and . . . an increase of the available CPU capacity on the disk servers (50% ?!). This will of course have an influence on the costing and sizing of the CERN computing facility.

The Storage Bits take

My system has 1 TB of data on it, so if the CERN numbers hold true for me I have 3 corrupt files. Not a big deal for most people today. But if the industry doesn't fix silent data corruption the problem will get worse. In "Rules of thumb in data engineering" the late Jim Gray posited that everything on disk today will be in main memory in 10 years.

If that empirical relationship holds, my PC in 2017 will have a 1 TB main memory and a 200 TB disk store. And about 500 corrupt files. At that point everyone will see data corruption and the vendors will have to do something.

So why not start fixing the problem now?

Comments welcome, of course.

Update: Peter Kelemen, one of the CERN researchers, kindly wrote in and pointed out that the it is the disks that are rated at 10^{14} , not the RAID card. There are no specs for the RAID cards. I've corrected it above.

