

Hard Disk MTBF: Where's the Reliable Reliability Data?

March 20, 2007

By [David Morgenstern](#)

The revelations in a couple of research papers on troubles concerning the MTBF specification for hard disk [reliability](#) sparked readers to suggest that there must be a better way to suss out a potential problem drive in the server closet. Furthermore, they have a good idea of who may have a finger on the real-world data and why that information isn't receiving an audience.

As I mentioned in a recent column on mean time between failure, a couple of papers presented at [FAST '07](#) (the USENIX conference on File and [Storage](#) Technologies) showed that annual disk replacement rates are much higher than predicted, the well-held belief in a burn-in phase for hard disk life cycle was wrong, and the SMART (self-monitoring, analysis and reporting technology) code in hard drives and storage management software—long touted by the industry as the best predictor of disk failure—was mostly a security blanket for IT managers.

Digging through a deep bucket of responses, many readers expressed some form of shrug about MTBF. Some may have believed in the rating in the past, but no longer.

"As you undoubtedly recall, MTBF was created by the drive manufacturers back when hard disks failed with much greater regularity. It was a way to reassure customers that manufacturers took reliability seriously, that drives were tested, and that by comparing these obviously inflated figures, you could assume that a drive with a million-hour rating was better than one with a half-million-hour rating," observed Barry Cohen, chief technical officer with technology analysis and consulting firm The Edison Group of New York.

"Actually believing that the drives themselves would last that long is a personal problem," he counseled.

Now, Cohen has a good point: MTBF is just a statistical measure. We're not supposed to believe it.

Still, storage managers do want to know what's going to happen to their hard disks. One side of our brain knows that each disk is just one example out of a production run in a product line that may have hundreds of thousands or even millions of units. The other side wants a date and time.

Do enterprise clients really need bigger and bigger hard disks? Maybe not. [Click here](#) to read more.

Ronald Major, manager at Sherwin Williams of Cleveland, said of course we all know that *all* drives fail and that MTBF has to be taken in the context of a population of drives. He believes that even storage vendors seem at times not to understand what MTBF means in the context of their own products.

"I asked a storage vendor about the MTBF of their drives, and he explained that they prefer to use 'mean time to [data loss](#).' I suppose that's supposed to make me feel better. But what does it really

mean?" Major said.

"You have a data loss on a Tier 1 array, and it sucks. I would be truly impressed with a vendor if they could tell me how many drives per year I can expect to fail, rather than how reliable their gear is. I would feel confident [then] that they know what they're talking about," he said.

Major hopes that, with the hubbub surrounding MTBF, perhaps the storage industry will present real-world reliability for their products, rather than some data from "idealized environments."

However, other readers said that the answer may be found by examining disks in controlled environments. In fact, the storage in question could be your own. But don't expect to know any more than you do now.

Former EMC employee Steve Smith, who runs an IT management consulting business in Bellevue, Wash., said that major suppliers of RAID, NAS and SAN to the enterprise and high-performance computing sites must have sufficient statistical information about MTBF.

"The controlled environment within an enterprise-class storage array is carefully monitored and controlled. The drives in these arrays are constantly compared by the suppliers. [But] these suppliers don't share their numbers with customers," Smith said.

Why not? He said suppliers believe customers would be shocked and react poorly.

"The simple fact is the internal story [the reliability statistics] doesn't match what customers assume. After years of letting their customers believe a fantasy, the suppliers are hesitant to reframe expectations around reliability," Smith continued.

"Why should the suppliers reset their customers' expectations about MTBF?" he asked. "Unless all of them do it simultaneously, someone will lose sales revenue. None of them will take that risk. And I wouldn't recommend they reveal the numbers without an offsetting benefit."

According to Smith, customers that purchase mass quantities of arrays could change this picture by demanding the real numbers from suppliers. But this tactic would take some backbone from the IT and purchasing departments.

"If sales revenue depends on showing the numbers, it will happen," he said. "The suppliers will ask for nondisclosure agreements. But if the large customer refuses to sign and says they will buy from another supplier who *will* reveal the numbers, supplier revelation is inevitable."

On the other hand, Marc Parpal Tamburini, a Hewlett-Packard product reliability engineer in Barcelona, Spain, suggested that false conclusions can be drawn by quick calculations and a lack of knowledge about statistics. He pointed to an interesting paper presented at [ARS 2005 \(International Applied Reliability Symposium\)](#) written by Sun Microsystems scientists David Trindade and Swami Nathan.

In "[Simple Plots for Monitoring Field Reliability](#)," the researchers discuss the problems with MTBF—statistical and customer-side—and recommend a "time-dependent reliability" model, which tracks a

customer's storage over time. By plotting a variety of data on the systems and their failures (and a bunch of other points) and then applying a number of statistical voodoo, customers can get a better picture of reliability.

One of the best methods to predict the failure of any device, storage or otherwise, is to simply count 30 days after its warranty. When the warranty is up, the product will fail. Or fall off your workbench onto the hard floor, warping the battery housing. Or a cup of coffee will be spilled on your desk and the liquid will drip down into the open vent and blow the power supply of the system stored below.

Such events rarely seem to happen under warranty.

In a similar vein, John Weinhoef, of Springfield, Ill., suggested that warranties can be used as a predictor of disk reliability. Now retired, he was the former manager of a 21TB high-performance computing storage operation.

"For enterprise operations, a better indicator was the maintenance rate charged for 24/7/365 service. The vendors knew what it was costing them to repair or replace failed units and adjusted their rates accordingly. When the projected maintenance cost over the next three years equaled new purchase cost plus a three-year 'warranty,' it was time to replace the disk subsystem."

According to Weinhoef, this meant replacing most disk subsystems every three years.

But when it comes to PC drives, he said that all bets are off as to reliability.

"The treatment in the field is ridiculous. The average person doesn't have a clue how delicate the drives are. I regularly see people ruining systems," he said.

He then related a story about dealing with a friend over an "Ethernet cable problem," or so it was described over the phone. It turned out that Weinhoef's friend had pushed the networking card completely out of the slot.

"When I got there she still had the system powered on, was slamming the box left and right about 24 degrees each way trying to shake the card back in place and was fishing around in the live box with an oversize, unbent paper clip. And people wonder why their systems fail," he said.

We can all smile at this and shake our heads knowingly. We would never, ever do anything as stupid as this in the enterprise or data center!

However, as I mentioned in my previous column, many current IT storage techs appear to have taken on a somewhat cavalier attitude toward the handling of drives in the field.

And I suggested that when folks toss around an iPod or a thumb drive or even some of the "ruggedized" external 2.5-inch notebook drives, they can pick up some bad habits when it comes to larger drives destined for desktops and servers.

Some of you thought I was being overly cautious.

Listen, I recall the same thing happening a generation ago with people in prepress shops handling Syquest cartridges and "real drives" housed in caddies for mirrored RAID systems. Both kinds of storage ended up being knocked around and given the same rough treatment.

Same difference nowadays and still no good for the data.

What do you think? Can your drives take a lickin'? Or do you baby your disks? Let us know [here](#).

Check out eWEEK.com's Storage Center for the latest news, reviews and analysis on enterprise and small business storage hardware and software.