

More Analysis on Disk Drive Mean Time Between Failure

A study by Carnegie Mellon University, where the researchers took roughly 100,000 [hard drives](#) from a variety of manufacturers, concluded that [hard drive](#) manufacturers were exaggerating their mean-time before failure (MTBF) ratings on hard drives – failure rates were, on average, 15 times the rated MTBFs. The drives were tested in various operating conditions, including real world scenarios. Researchers found that drive operating temperatures had little to no effect on failure rates: a cool hard drive survived no longer than one running hot. The types of drives used in the study ranged from Serial ATA drives, SCSI and even high-end fiber-channel (FC) drives. Carnegie researchers found that high-end drives did not outlast their mainstream counterparts: **“In our data sets, the replacement rates of SATA disks are not worse than the replacement rates of SCSI or FC disks. This may indicate that disk-independent factors, such as operating conditions, usage and environmental factors affect replacement rates more than component specific factors.”**

The study also found that the number one cause of drive failures was age. Drives tended to start showing signs of failure after roughly five to seven years of service, after which there was a significant increase in average failure rates (AFR). The failure rates of drives that were in their first year of service or shorter was just as high as those after the seven year mark. The average replacement rate of drives ranged from 2-13% annually, indicating that there is a need for manufacturers to re-evaluate the way a MTBF rating is generated. Worst of all, these rates were for drives with MTBF ratings between 1 million and 1.5 million hours. Carnegie researchers concluded that [backup](#) measures are a necessity with critically important data, no matter what kind of hard drive is being used.

Customers are replacing disk drives at rates far higher than those suggested by the estimated mean time between failure (MTBF) supplied by drive vendors, according to a [study of about 100,000 drives](#) conducted by Carnegie Mellon University.

The study, presented last month at the 5th USENIX Conference on File and Storage Technologies in San Jose, also shows no evidence that Fibre Channel (FC) drives are any more reliable than less expensive but slower performing Serial ATA (SATA) drives.

That surprising comparison of FC and SATA reliability could speed the trend away from FC to SATA drives for applications such as near-line storage and backup, where storage capacity and cost are more important than sheer performance, analysts said.

At the same conference, [another study of more than 100,000 drives](#) in data centers run by Google Inc. indicated that temperature seems to have little effect on drive reliability, even as vendors and customers struggle to keep temperature down in their tightly packed data centers. These results show how little information customers have to predict the reliability of disk drives in actual operating conditions and how to choose among various drive types.

Real world vs. data sheets

The Carnegie Mellon study examined large production systems, including high-performance computing sites and Internet services sites running SCSI, FC and SATA drives. The data sheets for those drives listed MTBF between 1 million to 1.5 million hours, which the study said should mean annual failure rates "of at most 0.88 percent." However, the study showed typical annual replacement rates of between 2 and 4 percent, "and up to 13 percent observed on some systems."

Garth Gibson, associate professor of computer science at Carnegie Mellon and co-author of the study, was careful to point out that the study didn't necessarily track actual drive failures, but cases in which a customer decided a drive had failed and needed replacement. He also said he has no vendor-specific failure information, and that his goal is not "choosing the best and the worst vendors" but to help them to improve drive design and testing.

He echoed storage vendors and analysts in pointing out that as many as half of the drives returned to vendors actually work fine and may have failed for any reason, such as a harsh environment at the customer site and intensive, random read/write operations that cause premature wear to the mechanical components in the drive.

Several drive vendors declined to be interviewed. "The conditions that surround true drive failures are complicated and require a detailed failure analysis to determine what the failure mechanisms were," said a spokesperson for Seagate Technology in Scotts Valley, Calif., in an e-mail. "It

is important to not only understand the kind of drive being used, but the system or environment in which it was placed and its workload."

"Regarding various reliability rate questions, it's difficult to provide generalities," said a spokesperson for Hitachi Global Storage Technologies in San Jose, in an e-mail. "We work with each of our customers on an individual basis within their specific environments, and the resulting data is confidential."

Ashish Nadkarni, a principal consultant at GlassHouse Technologies Inc., a storage services provider in Framingham, Mass., said he isn't surprised by the comparatively high replacement rates because of the difference between the "clean room" environment in which vendors test and the heat, dust, noise or vibrations in an actual data center.

He also said he has seen overall drive quality falling over time as the result of price competition in the industry. He urged customers to begin tracking disk drive records "and to make a big noise with the vendor" to force them to review their testing processes.

FC vs. SATA

While a general reputation for increased reliability (as well as higher performance) is one of the reasons FC drives cost as much as four times more per gigabyte than SATA, "We had no evidence that SATA drives are less reliable than the SCSI or Fibre Channel drives," said Gibson. "I am not suggesting the drive vendors misrepresented anything," he said, adding that other variables such as workloads or environmental conditions might account for the similar reliability finding.

Analyst Brian Garrett at the Enterprise Storage Group in Milford, Mass., said he's not surprised because "the things that can go wrong with a drive are mechanical -- moving parts, motors, spindles, read-write heads," and these components are usually the same whether they are used in a SCSI or SATA drive. The electronic circuits around the drive and the physical interface are different, but are much less prone to failure.

Vendors do perform higher levels of testing on FC than on SATA drives, he said, but according to the study that extra testing hasn't produced "a measurable difference" in reliability.

Such findings might spur some customers to, for example, buy more SATA drives to provide more backup or more parity drives in a RAID configuration to get the same level of data protection for a lower price. However, Garrett cautioned, SATA continues to be best suited for applications such as backup and archiving of fixed content (such as e-mail or medical imaging) that must be stored for long periods of time but accessed quickly when it is needed. FC will remain the "gold standard" for online applications such as transaction processing, he predicts.

Don't sweat the heat?

The Google study examined replacement rates of more than 100,000 serial and parallel ATA drives deployed in Google's own data centers. Similar to the CMU methodology, a drive was considered to have failed if it was replaced as part of a repair procedure (rather than as being upgraded to a larger drive).

Perhaps the most surprising finding was no strong correlation between higher operating temperatures and higher failure rates. "That doesn't mean there isn't one," said Luiz Barroso, an engineer at Google and co-author of the paper, but it does suggest "that temperature is only one of many factors affecting the disk lifetime."

Garrett said that rapid changes in temperature -- such as when a malfunctioning air conditioner is fixed after a hot weekend and rapidly cools the data center -- can also cause drive failures.

The Google study also found that no single parameter, or combination of parameters, produced by the SMART (Self-Monitoring Analysis and Reporting Technology) built into disk drives is actually a good predictor of drive failure.

The bottom line

For customers running anything smaller than the massive data centers operated by Google or a university data center, though, the results might make little difference in their day-to-day operations. For many customers, the price of replacement drives is built into their maintenance contracts, so their expected service life only becomes an issue when the equipment goes off warranty and the customer must decide whether to "try to eke out another year or two" before the drive fails, said Garrett.