

Understanding Reliability Metrics

Technology Paper

Guidelines for Determining the Most Suitable HDD for Your Application

The four criteria that typically take precedence in any hard drive (HDD) purchasing decision are: capacity, price, performance and reliability. Depending on whether you're a desktop enthusiast or an IT professional, each one likely gets a different priority. In general, enterprises expect unflagging uptime from HDDs that work around the clock, and they're willing to pay a premium for such reliability. On the other hand, mainstream PCs are more cost-sensitive and system builders know their customers expect lots of storage and responsiveness, so they balance those attributes against the price tag.

Of course, there's plenty of cross-segment experimentation going on. For example, might it be more cost efficient for a large data center to install inexpensive drives and simply swap them out as they fail? Or how about the power users willing to splurge on enterprise-class HDDs for a big, 10-bay tower?

But the most pressing question really is: How do you quantify HDD reliability and determine the most suitable drive for your particular application?

In business, total cost of ownership (TCO) matters most as it pertains to storage; it includes hardware procurement, power and cooling requirements, administration, maintenance, and eventual repair or replacement.

As you might expect, an HDD's reliability is one of the most influential variables in the TCO equation and the metrics used in defining reliability are persuasive. Customers compare them between storage vendors and among product families to guide purchasing decisions, while the manufacturers themselves depend on reliability metrics to shape warranty terms.

We are familiar with acronyms like mean time between failure (MTBF), annualized failure rate (AFR) and workload rate limit (WRL) used to predict how long a device might be expected to last, but their meanings are not the same. So which metrics matter the most, what do they mean, and how do they apply to your HDD's useful life?

Mean Time Between Failure (MTBF): Clarifying a Legacy Metric

The storage industry has used MTBF as the standard for quantifying failure rates for many years. Originally developed for the military, MTBF can be calculated in a variety of ways, unfortunately yielding different results. The specification is intended to quantify the probability that an HDD will fail based on an observed error rate over weeks and months of use.

MTBF, however, can be subject to misinterpretation. For example, an HDD might carry a 1M-hour MTBF, causing some customers to erroneously assume that the drive should be able to run 24 hours a day for a century or more. Obviously, this is theoretical.

MTBF is actually reached by dividing the total time a population of samples is under observation by the number of failures. If you have 1000 drives running for 1000 hours each, and five of them fail ($1000 \times 1000 / 5$), you've then derived a 200K-hour MTBF. Because MTBF is typically calculated during an HDD's useful life, consider a graph of failure rate over time, starting when a population of HDDs rolls off of the manufacturing line.

On the leftside of the graph (Figure 1), failures are shown to be more common as subpar components fail quicker than expected and drop out of commission early. (This is represented by the red dotted line in the graph below.) Over time, this levels off, yielding the useful life failure rate. It's much lower and affected by external factors like temperature and workload. Finally, as drives begin wearing out, the failure rate increases again, leaving you with a bathtub-shaped curve. To be clear, early and wear-out failures aren't part of the MTBF rating. What is calculated is the breadth at the bottom of the bathtub, as this relates directly to the HDD's warranty specs.

Rather than giving you any practical indication of how long an HDD should run trouble-free, MTBF better serves as a barometer of reliability. Using statistical analysis, the reliability of storage can be estimated over time. However, MTBF remains a difficult metric to use properly, in part because it is based

more on extrapolation than observation. That may be problematic for folks deciding which HDDs to buy, what types of workloads to run on them, and how long to keep them in service before eventual replacement.




Figure 1. A Conventional MTBF Bathtub Curve

Annualized Failure Rate (AFR): A More Accurate Measurement

Recently, Seagate began offering AFR specs on some drives in addition to conventional MTBF ratings. Although the two metrics are similar, AFR is a percentage estimate of the products that will likely fail due to a defect over a 1-year period. The subtle distinction lies in a switch from averages to percentages. MTBF is based on probabilities; AFR relies on HDD return numbers and analyses for its estimates. AFRs don't include HDDs returned without issues, those subjected to excessive shock or HDDs that were mishandled. Seagate expects AFR to help facilitate better service plans and spare unit strategies.

Remember: AFR, as with any reliability metric, makes specific assumptions about how many hours an HDD is running per year, endured motor start/stop cycles, and internal temperature, all of which affect reliability. But AFR doesn't necessarily help you get to the bottom of which HDD is right for your application. Enter the annualized workload rate (WRL).

Workload Rate Limit (WRL): Keeping HDDs Healthy

Storage devices keep track of power-on hours, lifetime reads, lifetime writes and more. Using that information, it's possible to calculate how much work an HDD is doing and whether the workload falls within a device's design parameters. If so, you can reasonably expect a flat AFR, whereas pushing a drive beyond expected parameters increases its probability of failure.

Annualized workload rate is the sum of lifetime reads and writes, multiplied by 8760 (the number of hours in a year) over total power-on hours. The result is expressed in TBs (terabyte) per year.

Seagate specifies a WRL—that is, the maximum read/write workload (also in terabytes per year) a given HDD is designed to

accommodate. Enterprise-class drives feature a WRL of 300TB to 550TB/year. Specialty-tuned drives, including NAS, archive, video and surveillance families, are rated for 180TB/year; Seagate's client portfolio (desktop and laptop HDDs) has a workload rate limit of 55TB/year, which is still a formidable amount of data within the client market segment.

It doesn't matter if you reach the WRL in the first month of use or over a year. As long as you're under the limit, the failure rate does not accelerate. You can run the SeaTools™ for Windows utility, which measures and displays the annualized workload rate of your drive. When the actual annualized workload rate exceeds a drive's WRL, reliability decreases.

Consider environmental conditions in your reliability projections. An insufficiently-cooled storage device, even operating under its WRL, will experience a higher failure rate than those at or under manufacturer specifications. Similarly, using a desktop HDD in a highly active chassis with significant vibration levels can also take a toll on the performance and reliability of the drive.

Total Cost of Ownership (TCO): Reliability Means Something

It is significant that so much of an HDD's reliability is in the hands of its operator. We're accustomed to seeing failures blamed on the mechanical nature of HDDs, but rarely ask whether a device was cooled properly or subjected to workload stresses outside of its design parameters.

Reliability metrics can be difficult to interpret. Ultimately they're really just variables in your TCO equation. Only a small number of businesses install so many drives that it's actually more cost efficient to simply run the least expensive desktop models outside of their specification and swap them out if they fail prematurely. But for everyone else, failures are expensive events. Choosing the best HDD suited to your workload, then carefully controlling its operating conditions, will help keep reliability consistent and costs down.

Conclusion

TCO is one of the primary metrics for nearly all data center operations. HDD reliability can negatively affect overall TCO if the drives used are not matched properly with the data center operating conditions. In addition to usage time and temperature, data center operators are strongly advised to consider the anticipated workload and its effect on reliability when selecting drives.

Seagate provides clear guidelines for proper selection of HDDs for any data center workload environment. Keeping HDDs at a low temperature within the specified range and within their usage time and workload specifications are necessary conditions for long-term drive reliability and improved TCO. Following these guidelines should ensure the best possible drive reliability and the lowest possible cost associated with HDD replacement, maintenance and testing.

seagate.com

AMERICAS
ASIA/PACIFIC
EUROPE, MIDDLE EAST AND AFRICA

Seagate Technology LLC 10200 South De Anza Boulevard, Cupertino, California 95014, United States, 408-658-1000
Seagate Singapore International Headquarters Pte. Ltd. 7000 Ang Mo Kio Avenue 5, Singapore 569877, 65-6485-3888
Seagate Technology SAS 16-18, rue du Dôme, 92100 Boulogne-Billancourt, France, 33 1-4186 10 00