TOSHIBA

MTTF, UER, MAMR

What's behind the most important HDD specifications?

Author: Rainer W. Kaese, Senior Manager
Business Development, Storage Products Division,
Toshiba Electronics Europe GmbH

HDDs deliver high storage capacities at low cost, but the various HDD models available on the market differ, sometimes quite considerably.
What is actually behind specifications like MTTF and AFR? What does the UER mean? And how are SMR and MAMR drives distinguished from each other? The following article provides an overview of the most important specifications and how they are interrelated to support users in selecting the models that best fit their specific applications.

The most important specifications of modern hard disks

Hard disk drives (HDDs) deliver high storage capacities at low cost, but the various HDD models differ considerably in some cases. Specifications such as mean time to failure (MTTF) and unrecoverable error rate (UER), as well as technologies employed like shingled magnetic recording (SMR) and microwave-assisted magnetic recording (MAMR), help us assess which drive is best suited to the respective application. This article provides an overview of the most important specifications and how they are interrelated.

The rising amounts of data in companies and private households are causing a rapidly increasing need for high-capacity storage media. Hard disks are still the most common form of data storage, because this classic storage media is being continuously developed and is significantly cheaper per unit of capacity than flash memory. Although flash memory prices are falling, they are doing so to a similar extent for HDDs. In the large online storage of cloud and corporate data centres, hard disks therefore remain just as popular as they do in network attached storage (NAS) systems and in storage solutions for video surveillance.

However, not every hard disk is suitable for every application – in fact, the drives are very specifically tailored to different areas of application. The manufacturers make these visible by dividing them into different HDD classes or product series. Usually they have drives for use in PCs, NAS systems, video surveillance solutions and data centres within their portfolio. These can differ in terms of interfaces, performance and reliability. The models for PCs and notebooks, for example, are usually designed for a daily operating time of 8-16 hours and an annual workload of 55 TBytes. They therefore have no place in NAS systems or servers – as, although they are cheaper, they would wear out quickly, due to the continuous use and the higher write and read loads. This would mean the probability of failure would increase.

NAS, surveillance and enterprise HDDs, on the other hand, must support 24/7 operation and cope with higher workloads. For NAS and surveillance drives it is generally 180 TBytes, while for enterprise drives it is 550 TBytes per year. The manufacturer's specifications for MTTF only apply if this value, known as the 'workload rate' or 'rated workload' in data sheets, is adhered to. If the rated workload is exceeded, the hard disk is not immediately expected to fail, but the probability of failure heightens and the MTTF decreases.

Reliable HDDs reduce maintenance effort

The MTTF is a statistical value that describes the average operating time until the HDD fails. Depending on the HDD model, it lies between 1 and 2.5 million hours. For a single drive this value has only limited significance – such a drive could fail at any time, so regular backups and RAID configurations are required to protect against a possible loss of data. With a larger number of hard drives, however, the MTTF helps to estimate how regularly failures could occur. With an MTTF of 1 million hours and 1 million drives, a drive failure each hour would be expected – or with 1,000 drives, a failure every 1,000 hours.

For hard disks in 24/7 operation, the annual failure rate (AFR) can be determined from the MTTF, which is more intuitive as a percentage value. Simplified, this can be calculated as follows: annual operating time of 8,760 hours divided by the MTTF in hours multiplied by 100. An enterprise drive with an MTTF of 1.4 million hours therefore has an AFR of 0.625%. In a data centre with 100,000 drives, 625 of these are expected to fail per year and would need to subsequently be replaced. If the operator opts for enterprise drives with an MTTF of 2.5 million hours, the AFR is 0.35% and only 350 hard disks are likely to fail in the same period – 275 fewer. Costly maintenance work is thus significantly reduced.

Suitable environmental conditions are essential

Less serious than failures are read errors, which are usually compensated for by internal error correction. This is not always successful, which is why the UER is a key hard disk figure. The sources of UERs can vary – with dust, electromagnetic radiation or a poorly executed write operation being among the main reasons why the stored data in a certain part of a disk may not be readable. For PC, NAS and surveillance HDDs it is 1 in 1,014 – with an un-correctable bit error occurring on average every 1,014 bits read, i.e., every 12.5 TBytes. In contrast, enterprise HDDs with a UER of 1 in 1,015 experience a read error only every 125 TBytes.

If the annual workload of a PC hard disk (55 TBytes) is distributed evenly, there is an un-correctable read error about every 2.7 months. If, however, it were confronted with the annual workload of an enterprise hard disk (550 TBytes), there would be an error every 0.27 months – about every eight days.

Both MTTF and UER are only achieved if the hard drives are used within the environmental conditions specified by the manufacturer. PC drives, for example, are usually designed to work at temperatures between 0 and 60°C, while enterprise drives are designed for 5 to 55°C operation – after all, they are installed in systems that are housed in air-conditioned rooms or data centres. The manufacturers also provide information on sensitivity to shocks and vibrations. NAS and enterprise HDDs are somewhat less sensitive in this respect than PC or surveillance HDDs, because several of them are incorporated into one device. Their rotational vibrations can amplify each other, which is why the NAS and enterprise models are equipped with special vibration sensors and control mechanisms that register and compensate for this.

Writing is faster on the outer data tracks

In addition to reliability, the most important factors for hard disks are performance and energy consumption. The highest performance is offered by special performance HDDs that work at 10,500 or 15,000 revolutions per minute (RPM) – but they have been increasingly displaced by solid-state drives (SSDs) for several years. In the meantime, enterprise HDDs with 7,200 RPM already deliver a sequential throughput of up to 280 MBytes/s and up to 400 input/output operations per second (IOPS). Storage systems with a few dozen of these drives achieve over 5 GBytes/s and more than 10,000 IOPS, which is sufficient for many modern applications.

However, the performance of HDDs decreases with their fill level because the outer, first-written data tracks on the rotating magnetic disks are longer and hold more data than the inner ones. During one rotation, the read-write head can simply write or read more data on the outside than on the inside. The 'sustained data rate' stated by the manufacturers in data sheets always refers to the outer tracks – further inside, the value can drop to about two-thirds.

Typically, all hard disk categories are connected with a SATA interface; only the performance HDDs mentioned above are exclusively available with a SAS interface. Today, the data rate of 6 Gbit/s (called SATA 3.3) is standard for the SATA interface, with backward compatibility to the previous versions with 3 and 1.5 Gbit/s.

Enterprise HDDs are available with either SATA or SAS interfaces, with SAS offering important features that SATA lacks – including higher signal strengths, end-to-end data protection and dual porting. However, SAS is expensive and has slightly higher power requirements. For companies that want to optimise their energy costs, however, there are other levers – first and foremost the modernisation of their hard disk infrastructure. Because most of the energy in an HDD is needed for the rotation of the spindles, the storage capacity and the workload have only a small influence on the power consumption. A few high-capacity HDDs are therefore more economical than many small HDDs.

The currently used SAS standard is called SAS-3.0, sometimes also called SAS3, and has a data speed of 12 GBit/s, also with backward compatibility to 6, 3 and 1.5 GBit/s. SAS4 with 24 GBit/s operation also exists, but is not used for hard drives because the data rate is limited by the mechanics.

Different block sizes offer flexibility

In the data sheets of enterprise HDDs, there is usually also an indication of the block size: 512n, 512e or 4Kn. This is the size of the logical blocks that can be written to or read from on a hard disk. In the past, this was always 512 Bytes, so the drives had native 512 Byte sectors – hence the 'n' in the designation. Later, larger sectors of 4 kBytes were introduced to write and read larger blocks, which facilitates the management of high-capacity hard disks. In addition, error correction also works more efficiently with larger blocks.

TOSHIBA

Modern file and operating systems can handle native 4 KByte sectors on hard disks, but older versions often cannot. Therefore, with 512e, a format was developed that relies on 4 KByte sectors, but emulates eight 512 Byte sectors in each of them. Older file and operating systems can write and read 512 Byte blocks as usual. When writing, however, there may be a loss of speed if the entire 4 KByte sector is not written. The hard disk must first read the entire 4 KByte sector to fill one or more of its emulated 512 Byte areas and then write the sector back – so an additional read operation is incurred.

The different block sizes of the enterprise models give businesses the flexibility to choose the drives that best fit their file and operating systems. Enterprise HDDs also offer flexible security options, usually referred to as self-encrypting drive (SED) and sanitise instant erase (SIE) in data sheets. SED is a hardware-based encryption directly through the hard disk, which is not only very secure, but also relieves the system in which the drive is installed. SIE is an option to securely erase all data immediately instead of through tedious overwriting.

SMR and MAMR provide more storage space

The various HDD models also differ in relation to the recording technology used: conventional magnetic recording (CMR), SMR and MAMR. CMR has been in use for years and was formerly called perpendicular magnetic recording (PMR) to distinguish it from the predecessor technology longitudinal magnetic recording (LMR). Since PMR has now been used for 15 years, today it is often referred to 'conventional'. However, PMR/CMR has reached its limit at 16 TBytes per drive. SMR increases storage density by working with overlapping data tracks, leading to higher recording density. Reading the tracks works as before, but when overwriting an existing track, the data of the overlapping track must first be read and then written back with the new data. This can cause fluctuations in write speed, but caches and caching algorithms will catch them. SMR is primarily used with PC and surveillance HDDs, because they do not have to handle sustained high write loads with random accesses. For occasional writes or sequential data streams, such as those delivered by surveillance cameras, SMR is ideal.

Higher-capacity enterprise HDDs, on the other hand, will rely on MAMR. A microwave-generating element on the write head helps to focus the magnetic flux, so that less magnetic energy is needed for writing. The write head can thus be smaller and write bits more densely. Currently, MAMR is used in 18 TByte and 20 TByte hard drives, and with advancements in this technology, hard drives up to 30 TByte can be expected in the future.

Since no data needs to be written via overlapping, MAMR is not subject to the restrictions and performance limitations associated with SMR technology. Hard disks with MAMR reach the same level as their PMR/CMR predecessors. A combination of MAMR and SMR is also technically possible, but this is currently practised to a very limited extent. With a combination of these two methods, it will not be long before the industry can benefit from capacities of up to 40 TBytes – but again, with SMR-typical performance limitations when it comes to random write access.

Summary

Capacity and price are not the only criteria that should determine the choice of a hard disk. Other specifications – such as reliability and ability to handle challenging operating conditions – must also be taken into account, so that applications are reliably supplied with data, and data stocks are not jeopardised. The manufacturer classifications of PC, NAS, surveillance or enterprise HDD are already a good guide. But if the user also understands the technical data, it is easier to find and optimally use the best possible hard disk for the respective application, because different applications require different hard disk models.

Toshiba Electronics Europe GmbH

Hansaallee 181
40549 Düsseldorf
Germany

info@toshiba-storage.com
toshiba-storage.com

Issued 10/2022