



open-e

Open-E JovianDSS-Powered

**Data Storage Appliances
with Toshiba HDDs**

WWW.OPEN-E.COM

Introduction

Technologies behind hard disk drives (HDDs) have been well tested and are currently well known, as this type of disk has been present on the market for more than 50 years. Throughout this long period of time HDD technologies have been steadily developing, increasing the capacity, decreasing the physical size, and becoming more sustainable and affordable.

The Power of HDD Data Storage Appliances

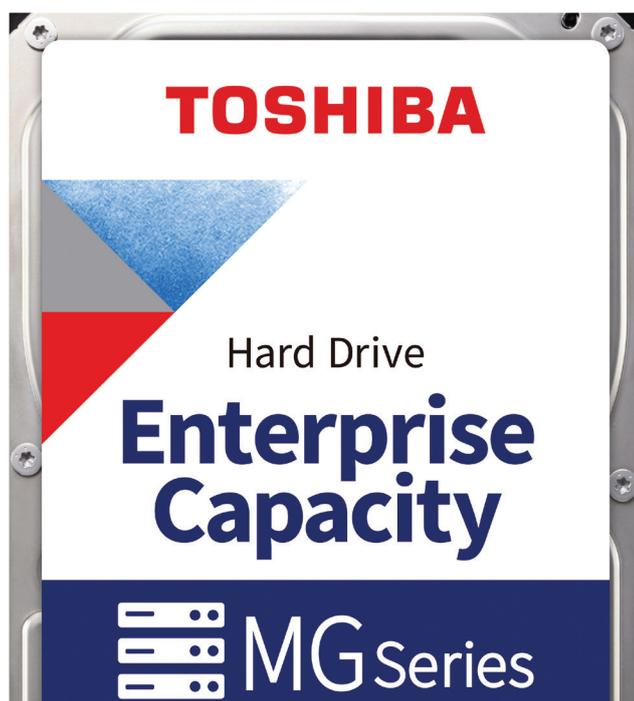
In large storage appliances (dozens of TB to PB range) HDDs have been and will still be the working horse for the base data storage. Why? HDDs simply offer the highest storage capacity at the lowest cost. Components based on alternative data storage technologies are used as well, but have either order-of-magnitude higher cost per capacity (Flash/SSD) or are not suitable for online data.

Modern Enterprise HDDs offer:

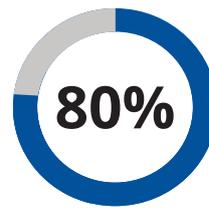
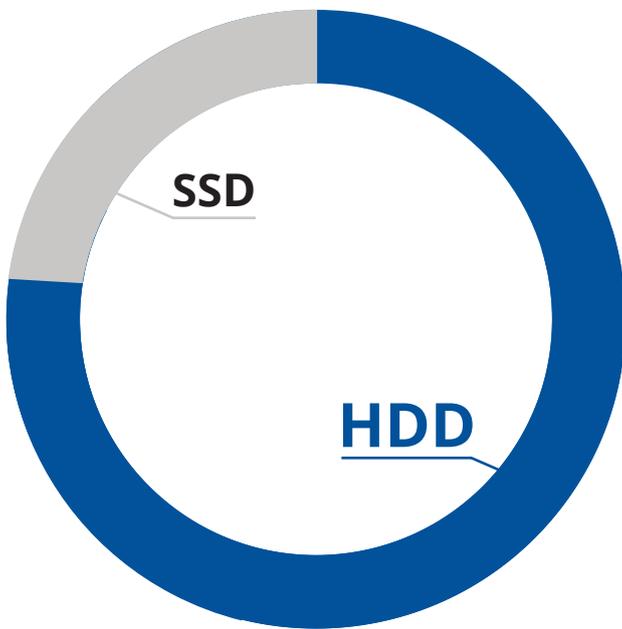
- Up to **20TB in the typical 3.5" form-factor**
- The workload of up to **550TB/year and 24/7-Operation**
- Yearly failure rates as low as 0.35% (equivalent to a Mean-Time-to-Failure MTTF of 2.5M hours)
- **5 years of warranty** time
- Out of 1000 drives only 3-4 parts to be expected to fail within a year of demanding enterprise operation
- **Infinite overwriting options** without the risk of damaging the surface (HDDs are not subject to endurance limitations like SSDs where the flash cells tend to "wear out" at over-writing/erasing)

Even if alternative technologies were available at comparable or competitive cost and reliability, still the pure amount of data to be stored today can only be served by HDDs. **In 2021, the entire storage component industry manufactured approximately 1700EB (that is 1,700,000,000TB) of total capacity.**

Only about 340EB of this was based on SSD technology, with the 1360EB majority being provided through HDDs. Within the next years and the explosion of data capacity to be stored there will be no chance to replace HDDs by other technologies simply due to the total amount of data storage required.

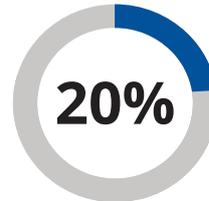


Total of 1700EB (1,700,000,000TB) of storage capacity in 2021



HDD

The majority of 1360EB was provided through HDDs



SSD

Only about 340EB of this was based on SSD technology

Capacity

The regular capacity of typical storage components may be divided into several groups:



16GB, 32GB and 64GB. This range is typically found in older and smaller devices. Modern HDDs do not scale down to this range, so flash memory-based components are used.



120GB and 256GB. Also for this range, HDDs do not scale down. Also, as SSDs are much faster, this capacity range is today covered by SSD.



500GB, 1TB and 2TB. This is the entry-level capacity for HDDs. HDD is still significantly cheaper than SSD (but slower). So for cheap entry level devices HDD is used, for professional laptops, desktops and workstations SSDs are used at significantly higher costs for the storage component.



4TB, 6TB, 8TB. This range is used in media-centric devices, workstations and video surveillance recorders as internal storage device, but also in enterprise servers and network attached storage appliances, typically as a set of few HDDs in RAID configuration for redundancy in case of drive failures. SSDs in this capacity range are very expensive and used only for very high-performance solutions.



10 to 20TB. This is the domain of enterprise and cloud storage servers.

However, the capacity numbers are not the only numbers that should be considered when choosing an HDD drive. One should also pay attention to the **TB/year workload rating**, which illustrates how much capacity may be processed in a more extended period of time. For example, the aforementioned **Hard Drives from Toshiba have a workload rating of 550TB/year.**

Lifespan

There are numerous debates being held about whether SSDs are going to reach the lifespan of HDDs or not. However, although the supplier warranty period is 3-5 years for an HDD, latest numbers have shown, that it can reach up to 10 years if the specific conditions are followed. On the contrary, SSDs are claimed to have a 5-year lifespan under optimal operating conditions.

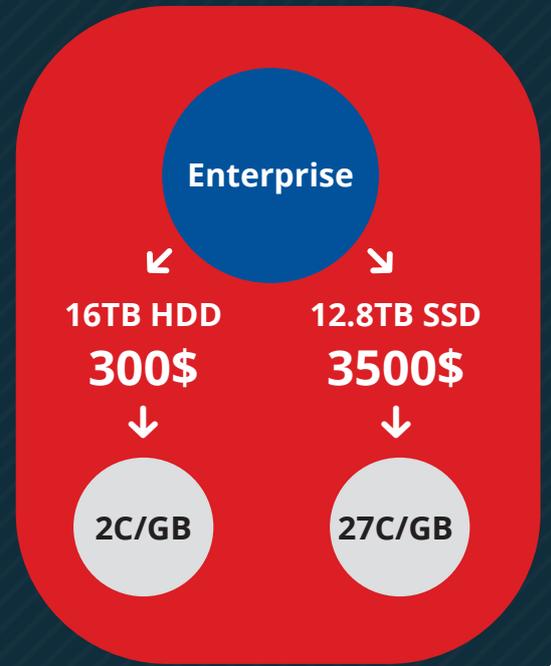
Total Cost of Ownership

The benefits of a hard disk drives are that they are a proven technology, and are frequently less expensive than a solid state drives for the same amount of storage. Currently, HDDs are also available with more storage space than SSDs.

SSDs are significantly more expensive than hard drives in terms of dollar per gigabyte.

The differences are even more visible when taking into consideration high-capacity 3.5-inch hard drives. For instance, **a 16TB 3.5-inch enterprise hard drive for about \$300 can turn the cost per gigabyte to less than 2 cents**, while comparable enterprise SSDs are usually ways above even 20 cents per GB.

Also, due to the fact that HDDs are based on a more established technology it is more likely they will remain inexpensive.

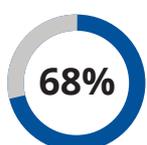


Sustainability

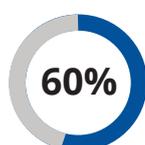
Hard Disk Drives, strictly speaking - their magnets, consist of imported rare elements. Due to the rapid data storage solutions market expansion, the demand for such materials increases as well, which in turn translates into very probable limits when it comes to the accessibility of those critical materials. **It has been estimated that approximately up to 70 million(!) hard disk drives reach the end of life every year only in the United States.** This enormous number shows not only the demand for hardware but most importantly in the case, the extent up to which hardware waste has begun to grow.

Luckily, HDDs can be reused and recycled. This is an exceptional feature of this kind of hardware, as hardware recycling helps preserve valuable materials and therefore reduces the negative impact on the environment. It has been researched that about **68% of materials that build HDD will be recovered** from recycled drives. This means that the overall weight of recycled elements in 2050 would be of 4.73 thousand metric tons with about 1.5 million metric tons of CO2 emissions would be avoided. In the best scenario (with an assumption of the highest material recovery) the **overall weight of recycled elements in 2050 would increase to more than 100,000 metric tons and 2 million metric tons of CO2 emissions to be avoided.**

HDD facts:



Recycled HDD materials



Old hard disk drives still in use

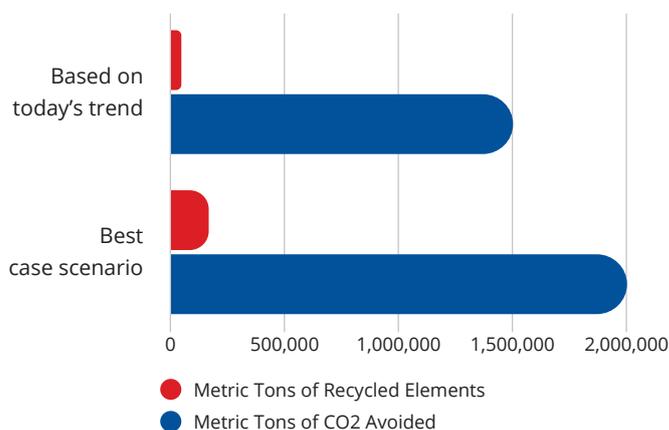


Existing magnets can be regained unaffected

Still, HDD reuse is dependent mainly on the trust that users have for the data-wiping technologies, but the fact is that data-wiped drives indicate the hardware has the data lastingly erased and any attempt to recover the wiped data will fail. It has been established that even if recycling is considered to be the ultimate end of an HDDs' lifecycle, **in fact reusing old drives is beneficial to the economy in the longer run.**

All in all, whether users decide to reuse or recycle, it is still clear proof of HDDs' environmental advantage.

2050 Forecast:



Source: <https://www.nrel.gov/news/program/2021/in-a-circular-economy-hard-drives-could-have-multiple-lives-in-the-future.html>

Comparison of HDD and SSD

Factor	HDD	SSD
Mechanism	Magnetic rotating platters	NAND Flash
Size	Large	Compact
Industry Standard Dimensions and Interface	Included	Included
Read/Write Speed	250 MB/s	500-4000 MB/s
Encryption	Full Disk Encryption as option	Full Disk Encryption as option
Capacity	Up to 20TB	Up to 30TB
Cost per GB*	4-6 cents (Low cap client) 2 cents (high cap enterprise)	10 cents (low cap client) >20 cents (high cap enterprise)
Sustainability	Can be recycled or reused but require more power usage	The manufacturing process has a big impact on the environment and cannot be reused
CO2 emission	1TB HDD = 99-199kg of CO2/ 5-10years	1TB SSD = 184-369kg of CO2/ 5-10 years
Noise and vibrations	Spinning, clicks, and vibrations	No vibrations and no sounds
CPU Power (I/O wait)	7%	1%
Failure Rate	1.38%	1.05%
Lifespan	Higher	Lower

*internal 2.5-inch hard drive

HDDs are magnetic carriers with movable head that moves over circular magnetic plates. SSDs are in turn classified as flash memories, made only of NAND flash chips. The semiconductor design makes reading and writing files on the SSD faster, because the SSD has no moving parts that could delay its operation. Also, the moving parts of the HDD are responsible for making sounds during operation and rotation. At the same time, SSDs without these moving parts work silently.

But SSD media also have a number of disadvantages. First of all, they have a shorter lifespan due to the limited number of read and write operations, they are much more expensive (poor price/performance ratio) and, last but not least, they cannot be easily recycled (not even mentioning the fact that it is impossible to reuse them), so the **environmental impact of an SSD is also much bigger** than in the case of an HDD.

As for HDDs - the price aspect returns here, as **big capacities can be bought inexpensively.** In addition, in the event of a failure, the data from the HDD can still be read, while an SSD usually loses its memory irretrievably, and their recovery is unfortunately time-consuming and difficult.



Performance Challenges - HDD vs SSD

Of course, a single HDD is slow (about 280MB/s and 150 IOPS reading / 500 IOPS writing). Compared to an SSD with several GB/s of sequential speed and 100.000 of IOPS, the single HDD could only be considered as a slow secondary/archive component. And in fact, the majority of "single component" systems such as PCs, Workstations etc. are nowadays using SSDs as the system's storage components.

However, for enterprise and cloud storage appliances, **it is not about the performance of a single component.** Due to the requirement for the highest capacity, also many HDDs are required. Typical solutions, for instance build of a ZFS-based Open-E's JovianDSS, range from a **practical minimum of around 8 HDDs up to arrays of hundreds of drives in JBOD expansion chassis.** And these high numbers of (slow) HDDs combined in the right way add up to a total performance of many GB/s range of sequential access speed at some thousands of IOPS. Not as high as used from SSD or All-Flash appliances, but still very suitable for storing large online data.

And these large and agile base **arrays of HDDs can still be accelerated with read and write cache structures based on lower capacity modules of Flash/SSD technology.**

The Open-E JovianDSS Caching Approaches

To overcome the performance limitations of data storage based on HDD drives several types of Open-E JovianDSS caching approaches can be deployed so as to help in accelerating both the read and write operations.

Read Cache in Open-E JovianDSS

A read cache uses a fast memory, located closer to a processor core to store copies of the frequently or recently read data (in the Open-E JovianDSS it's a combination of both). When the application needs to read the data, it shows **spectacular accessibility and speed**, since in the majority of cases it can be read directly from the fast memory unit instead of slower HDD data drives.

The regular read cache mechanism, called the 1-level cache (or ARC in ZFS), is placed in RAM, which makes it an **extremely fast and effective**, but also a bit costly solution. It's also not easily scalable due to the TCO and business requirements of a balanced budget.

Luckily **Open-E JovianDSS is based on ZFS**, so there is also a **2-level read cache available (L2ARC in ZFS)**. It is an extension of the RAM-based cache used when cached data is about to be removed from it. L2ARC is usually located on a drive that is not as fast as RAM, but is still well-performing (i.e, SSD, or NVMe), so the data can be easily read. This way, you **save the RAM resources, but still have fast access to the cached data**.

Write Log in Open-E JovianDSS

Write operations can be supported in a similar manner. It can be achieved by means of an **additional write log device (SLOG) which works as a buffer** where all the data is written before eventually being stored on HDDs. This way **very good write operations performance can be achieved** as the data is written on a fast SSD or NVMe write log drive first, and later on the final HDD data disk, while the second part of this operation doesn't engage the RAM anymore.

But the write log provides a few more advantages. For instance, it **mitigates the RAM loads in terms of write operations** since the data is promptly transferred from the RAM to the write log, without hangs caused by the writes queue on an HDD drive. Moreover, since RAM is a volatile memory and loses the data once the power is off, using the write log (non-volatile memory) is also an extra way to protect your data in case of power shortages or power supply failures.



Toshiba Enterprise HDD Models

Toshiba offers 3.5" Enterprise Capacity HDDs for various capacity points. From currently 18TB, and soon also 20 and 22TB all the way down to 1TB for legacy and embedded/industrial applications. 18TB is based on the 9th generation of Toshiba HDD technology (MG09-Series), and 20TB will be based on the 10th Gen MG10. The 8th Generation MG08 offers 16TB models, but also optimized versions for lower capacities of 4/6 and 8TB. The 7th gen MG07 covers 12 and 14TB, 6th gen 10TB and 1 and 2TB are still being manufactured based on the 4th generation MG04. With this wide lineup still being supported, Toshiba customers and users of Toshiba hard disk drives can maintain their storage infrastructure investments for a long time with optional expansions and drive replacements, at the same time not having to worry about termination of the used HDD models.

All drives are available with 6GB/s SATA or 12B/s SAS interface. SATA drives are used for simple configurations like a bunch of locally attached drives within a server chassis. The SAS interface offers two independent data channels and a higher data/signal protection, so SAS drives are the choice for systems with high availability (HA) requirements. But SAS drives with their interface features offer also advantages for larger and complex configurations like usage of external enclosures with cabling. The dual path to the hard disk drives enabled by SAS interface provides redundancy against failure of cables and expander modules.

TOSHIBA

open-e

Toshiba is a world leader and an innovator in pioneering high technology. It is a diversified manufacturer and a marketer of advanced electronic and electrical products spanning from information & communications systems to digital consumer products. It's also involved in electronic devices and components, power systems, industrial and social infrastructure systems as well as home appliances. Toshiba Electronics Europe GmbH, Storage Products Division, in Düsseldorf, Germany, markets hard disk drives with a particular focus on consumers and retailers in Europe. For over 50 years, Toshiba has been developing and manufacturing storage solutions used by most major IT and consumer electronics brands. For more information, visit www.toshiba-storage.com

Open-E, founded in 1998, is a well-established developer of IP-based storage management software. Its flagship product, Open-E JovianDSS, is a robust, award-winning storage application that offers excellent compatibility with industry standards. It's also the easiest to use and manage. Additionally, it is one of the most stable solutions on the market and an undisputed price performance leader. Thanks to its reputation, experience, and business reliability, Open-E has become the technology partner of choice for industry-leading IT companies. Open-E accounts for over 38,000 installations worldwide. Open-E has also received numerous industry awards and recognition for its product, Open-E DSS V7. For further information about Open-E, its products and partners, visit www.open-e.com