Comparing SSD and HDD Endurance in the Age of QLC SSDs

SSDs brought unique benefits into client and enterprise storage — high speed, low power and low latency to drive emerging (and standard) applications toward new performance thresholds. SSDs also introduced a new concept into the storage market: storage devices wear as they are written and require a warranted write endurance limit (SSD endurance) as a known variable versus HDD time-based warranties. Different types and sizes of SSDs offered different wear ratings, but the idea of storage devices wearing predictably as they were used was new. In the early days of SSD adoption (around 2007), HDDs didn't have warranted endurance ratings; today, ratings are much more broadly published.

The market saw initial user trepidation, partly because users didn't know their I/O patterns and had difficulty estimating them. As a result, many storage architects deployed SSDs with endurance ratings far greater than their workloads needed.

In this paper we will discuss how SSDs wear versus HDDs and how they can be compared based on workload/application needs, as well as:

- Developments in HDD endurance statements
- Converting device endurance ratings into common units
- Two examples of HDDs with wear ratings (workload limit ratings)
- How capacity-focused, enterprise SSD and HDD endurance differs
- How the introduction of new QLC NAND has created an endurance paradigm shift

Understanding how SSDs and HDDs are affected by different application environments will enable you to make purchase decisions based on your workload needs.

Note: This paper is focused on warranty statements related to endurance. These statements are often found in product manuals or data sheets. It does not contemplate actions taken when those values are reached. Although current at the time of publication, values used in this paper are subject to change. Your results may differ from those stated herein. This paper uses the terms "warranted endurance" and "endurance" interchangeably.

Storage Wear –
A Printing Analogy

Rated or warranted endurance wasn't a unique concept in the computing ecosystem. Printer cartridges are an example of computing devices that wear. The two are (somewhat) analogous in this way.

We have long accepted that printer cartridges can print different numbers of pages depending on how much print is on the page. If we print pages that have a lot of images and require a substantial amount of ink, we accept that the cartridge will have a shorter useful lifespan versus if we print pages with a little text. The less we print (or effectively wear the cartridge), the longer the cartridge will last.

SSDs and HDDs (with wear limit ratings) are like print cartridges in this way: they wear based on use. HDDs wear based on data written or read, while SSDs wear based on writes.

SSDs and HDDs incur wear differently.

SSDs wear when written (expressed as drive writes per day or DWPD).

HDDs with a workload limit rating incur wear when read *or* written and are compared on a DRWPD basis (drive reads or writes per day, expressed as a ratio of the drive's usable capacity).

If you make drive purchase decisions based on rated endurance, we show that for many read-focused workloads, low endurance SSDs meet or exceed the DRWPD of some capacity-focused enterprise HDDs.

SSD Wear

The NAND storage used in SSDs is different from the magnetic storage used in HDDs. The magnetic storage on HDDs supports write in place; if there is already data in the physical location to be written, the existing data can be directly overwritten with new data.

This is a single-step process. However, when data is present in NAND (even data that has been marked as no longer in use), NAND must be erased before it can be written (programmed). This two-step process is called a program/erase (P/E) cycle, and SSD endurance is a function of the number of P/E cycles for which the NAND is rated. (Additional details on this process may be found in this Micron Brief.)

NAND wears when it is written (this becomes important in the following sections), but only incurs negligible wear when read.

SSD Capability to Absorb Wear

As SSD adoption expanded, the industry grew accustomed to the idea that storage devices had a rated wear value. SSDs usually refer to this rating as drive writes per day, or DWPD (or less commonly, as total bytes written, or TBW). These values are related: DWPD = TBW / [(drive capacity) * (warranty period in days)]. We will use DWPD when discussing SSD warranted endurance.

Because the idea of storage incurring wear when written was new, many system designers overestimated the amount of DWPD they needed from SSDs. In the early days of SSDs, 10, 20 or more DWPD was normal. It ensured a safety margin. It was also due to drive capacities being significantly smaller (most early SSDs held less than 1TB). However, trends showed DWPD rapidly decreasing as drive capacities increased and the industry gained a better understanding of workload read/write profiles.

Figure 1 shows how DWPD requirements have decreased over time.

The type of NAND flash technology that SSDs are built on differs in several ways. The most fundamental difference among NAND types is the number of bits stored in each cell. The number of bits (0s or 1s) is controlled by the number of charge states for which the NAND is designed.

At one end of the spectrum, single-level cell (SLC) NAND holds one bit per cell. At the other end of the spectrum, new quad-level cell (QLC) NAND supports four bits per cell.

These different types of NAND (SLC, MLC, TLC and QLC) are also typically rated to support a different number of P/E cycles. (As any NAND approaches its rated P/E cycle count, it becomes increasingly difficult to determine if the data stored is a 0 or a 1.) There are some techniques that can extend the number of P/E cycles supported, but at some point, the NAND is worn to the point where it is retired. The state of wear for the SSD is managed internally and is easily monitored by the OS or storage system, making the lifespan of a drive very predictable.

Generally speaking, as the number of bits per cell increases, the number of P/E cycles for which the NAND is rated decreases. This is due in part to the additional complexity introduced by more charge states. The effect is reduced warranted write endurance of storage devices using that NAND, along with reduced cost per bit stored, as illustrated in Figure 2.

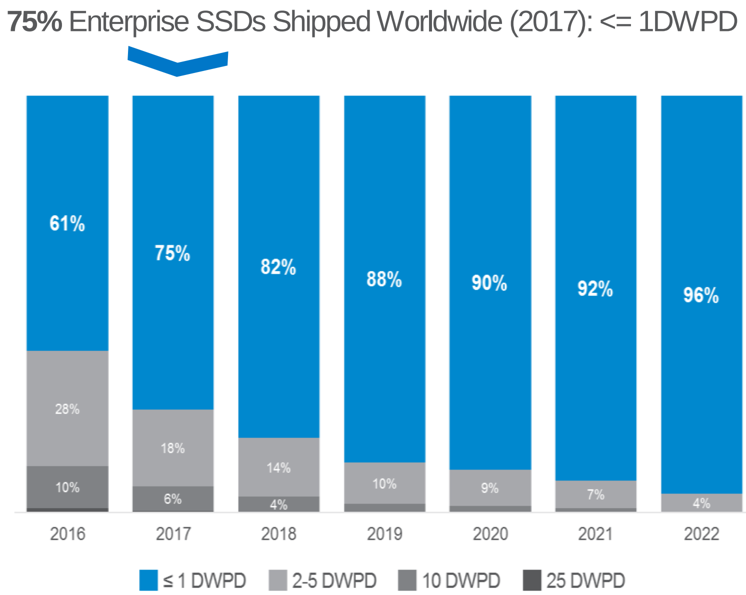

Figure 1: SSD Endurance Trends
(Source: Analyst Consensus, Forward Insights Datacenter, May 2018)

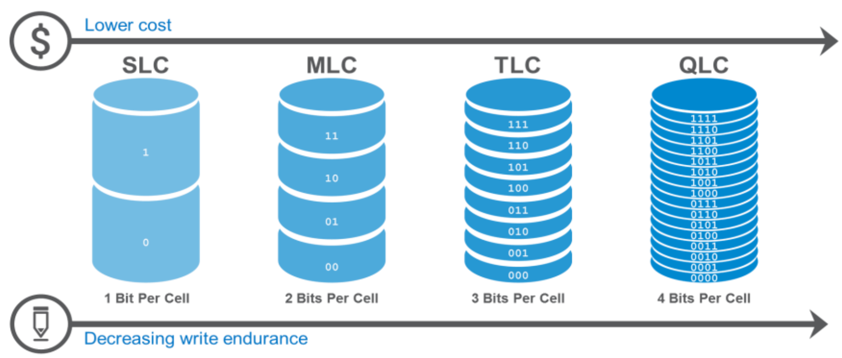

Figure 2: Cost, Complexity and Write Endurance Trends

Changing Trends in SSD Endurance Requirements

Advances in NAND technology spurred the enterprise SSD market to broaden adoption and at the same time rethink endurance requirements and re-analyze workloads and the amount of data they write. There was a general realization that many workloads wrote far less than initially anticipated. This drove broad reconsideration of how much write endurance was really needed. Given the reduced price per gigabyte, combined with increased per-drive capacity introduced by advances in NAND storage technology, more workloads could take advantage of SSDs. Workloads that wrote less and benefited from flash (like a read cache) became prime candidates for denser, lower-warranted endurance NAND.

This culminated in the introduction of QLC NAND and Micron's release of the first QLC SSD to the market in 2018.

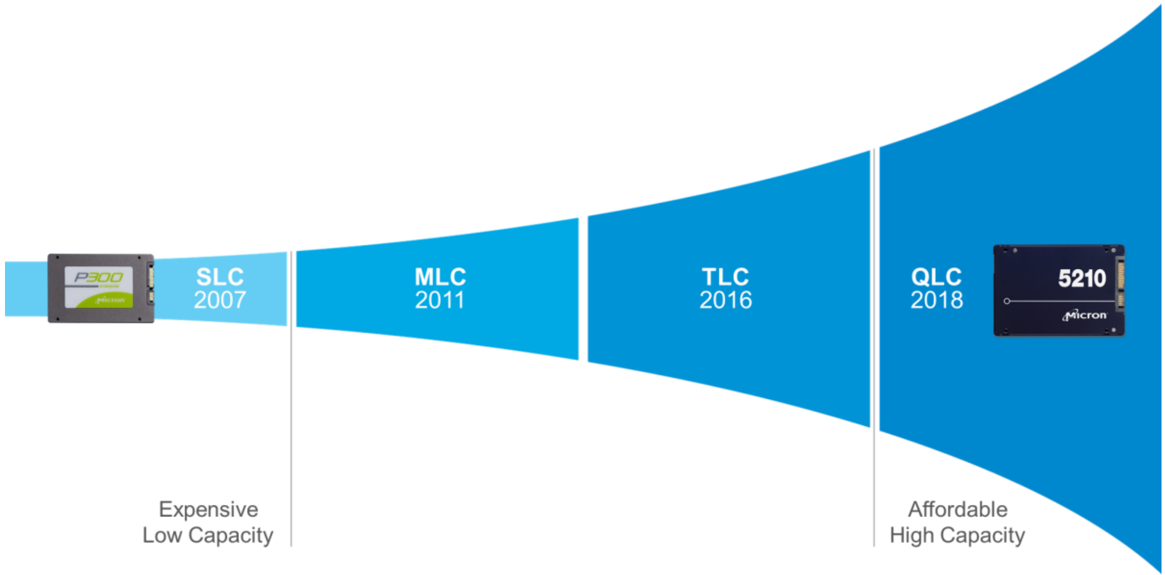

Figure 3: SSD Trends

HDDs Adopt Workload Ratings

HDD capacity has also grown in the same timeframe (using different technologies than SSDs). At the time of this paper's publication, leading HDD vendors were offering 14TB 3.5-inch form factor products. As HDD capacity grew, the breadth of HDD types also grew, with some designed for specific use cases. As their capacities increased, some HDD designs began to adopt a "workload limit" rating as part of their standard specifications.

Workload Limit Ratings

HDD workload limit ratings are very different from SSD DWPD ratings. Both are warranted values (exact warranty terms will differ between brands, models and technologies), but while SSD warranties typically state endurance in terms of the amount of data written, HDD warranties typically state workload limits in terms of bytes ***written and/or read.*** (The reasons for HDDs adopting specific workload limit ratings are beyond the scope of this paper.)

This means that different I/O types wear SSDs and HDDs differently, as shown below.

Drive Type	Reading Data	Writing Data	Varies w/Pattern
SSD	Minimal/no wear	Causes wear	Yes
HDD	Causes wear	Causes wear	No

Table 1: Drive Wear Incurred by I/O Type

Comparing Endurance: SSD TBW and HDD Workload Limit

In this section we compare the warranted endurance of two typical, enterprise-class HDDs designed for high-capacity storage.

Each example is an 8TB enterprise-grade HDD, one designed for general purpose storage and the other as a bulk storage device (active archive). These HDDs are only examples; exact workload limits may vary by manufacturer, model, capacity, generation or many other factors. Some models don't have a specified workload limit; therefore, your results may vary.

HDD Workload Limits

To illustrate HDD endurance, we'll examine two enterprise-class 7200 RPM drives. Note that other HDD types may have different endurance ratings. Table 2 shows their relative values and data sheet workload limit ratings. (Workload limit ratings are typically expressed in TB/year.) Since HDD workload limit ratings include both read and write I/Os, we can express a workload limit in SSD-familiar terms, or DRWPD. The DRWPD value in the table below was derived from each drive's data sheet workload rating using simple math.

HDD	Workload Limit Rating (TB Read or Written per Year)	DRWPD	HDD Capacity	Spindle RPM	Warranty
Example 1	550TB	0.19	8TB	7200	5 Years
Example 2	180TB	0.06	8TB	7200	3 Years

Table 2: Workload Limit Ratings by Example HDD

The DRWPD values noted in Table 2 were calculated as follows:

Example 1 HDD: 8TB, 3.5-inch 7200 RPM general purpose enterprise storage

Workload limit rating: 550TB per year
Warranty duration: 5 years
Drive capacity: 8TB

Warranted drive reads or writes per day: (550TB per year/365 days per year)/8TB capacity = 0.19

Example 2 HDD: 8TB, 3.5-inch 7200 RPM active archive enterprise storage

Workload limit rating: 180TB per year
Warranty duration: 3 years
Drive capacity: 8TB

Warranted drive reads or writes per day: (180TB per year/365 days per year)/8TB capacity = 0.06

DRWPD has not historically been a frequently used industry term, but in the age of QLC SSDs, it will likely become a more important point of comparison, as SSDs increasingly compete with HDDs for primary data storage. We create and use the term here to enable clearer, more direct comparisons with SSD DWPD ratings.

We now can graph HDD workload limits as DRWPD, as shown in Figure 4.

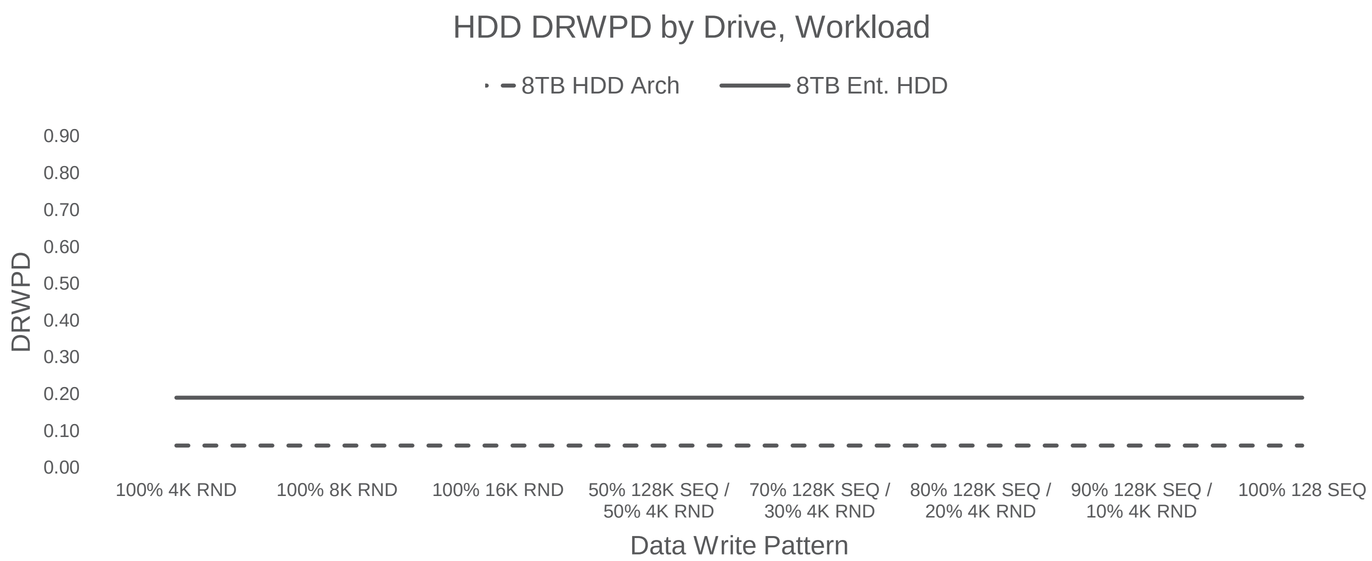

Figure 4: HDD DRWPD by Write Pattern (workload)

Note that DWPD is constant for these drives. This is expected since rated endurance is also fixed (workload-independent). The x-axis shows several different I/O mixes, or data write patterns. Figure 4 indicates what percent of all write I/O traffic is random versus sequential (and common transfer sizes for those accesses). Figure 4 also assumes that the HDD can generate sufficient I/O to realize these DRWPD values. This may not be the case with all HDDs.

Comparing SSD DWPD and HDD DRWPD

Earlier this paper noted that NAND wears when it is written, not when it is read. This means that NAND-based SSDs also wear when written, not when read. (Note that read disturb, a phenomenon of SSDs, incurs a slight amount of wear, but it is negligible.) SSDs incur more wear when I/Os are small and randomly placed. The opposite is also true — SSDs incur less wear when the I/Os are large and sequentially placed. (See this Micron Brief for additional details.)

Prior to the introduction of QLC NAND technology, SSDs were rated at a fixed DWPD value. That means that their DWPD ratings did not change with applied workload. SSD DWPD was typically specified with a small (4K), 100% random write workload.

The difference between SSD wear applied by different write patterns and SSD DWPD ratings should be understood: while different write patterns may cause more wear, prior to the introduction of QLC NAND, SSD-rated DWPD did not change.

When QLC technology was first shipped into the enterprise SSD market (the Micron 5210 ION was the first such SSD shipped), QLC NAND wear characteristics and workload understanding had matured. With that greater understanding came a new way of expressing QLC-based SSD endurance, and the data sheet for the Micron 5210 ION SSD gave workload-specific endurance expectations — an industry first that was quickly followed and adopted by other manufacturers of QLC SSDs.

This means that the rated DWPD on the Micron 5210 ION SSD varies based on the type of write I/O — small and random I/O means lower DWPD; a heavier mix of large and sequential I/O means higher DWPD.

Figure 5 shows the DWPD of a 7.68TB Micron 5200 ECO SSD (whose DWPD is noted as a constant in the data sheet) and a 7.68TB Micron 5210 ION SSD (whose DWPD varies with workload), both overlaid on the HDD DRWPD that was previously noted and discussed in Figure 4.

Note that Figure 5 uses "DxPD" to indicate that it shows both DWPD and DRWPD. When DxPD references an HDD, it is understood to mean DRWPD; when referencing an SSD, it is understood to mean DWPD.

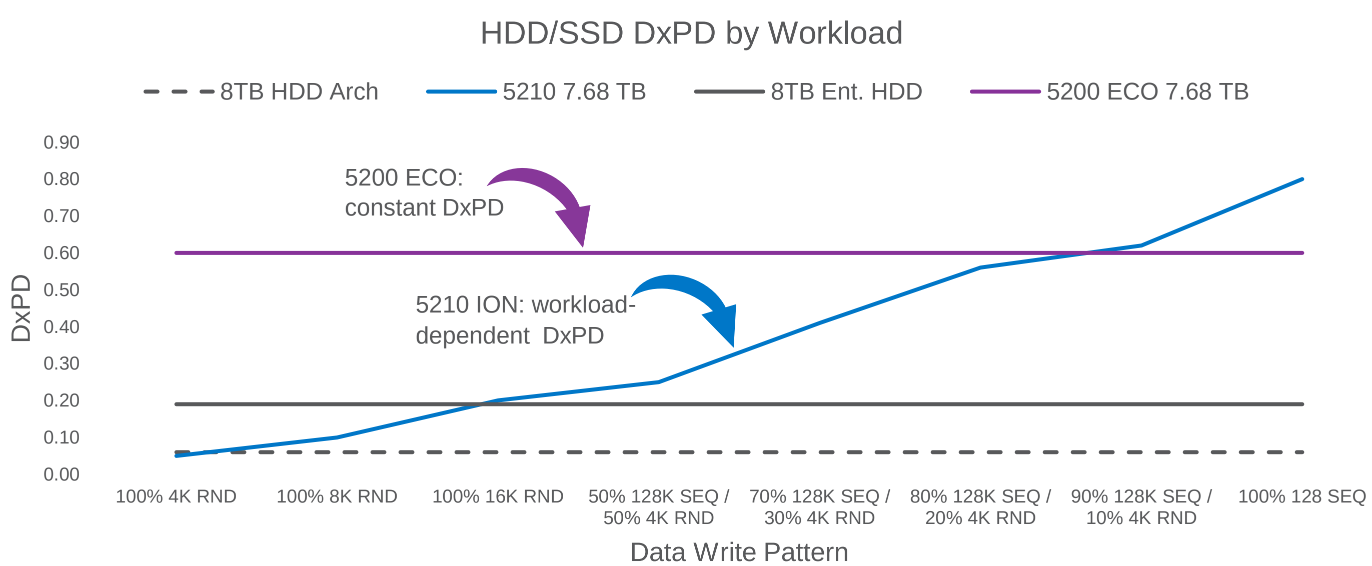

Figure 5: Micron Read-Intensive Enterprise SSD DWPD vs. HDD DRWPD (DxPD) by Write Pattern (workload)

Figure 5 compares only enterprise SATA drives (SSD and HDD) to ensure a fair comparison. The 5210 ION SSD DWPD ratings trending up and to the right indicate that as the data write pattern changes from small random I/O (at left), to larger, sequential I/O (right), DWPD increases on the 5210.

The 5200 ECO and HDD DWPD data sheet ratings do not follow this trend. Both are independent of the I/O type with small random and large sequential, resulting in the same DxPD (appearing as a fixed, horizontal line).

Figure 5 also shows:

- The 5210 and archive-class HDD have similar DxPD if all write traffic is small (4K) and random. When the I/O size reaches 8KB, the 5210 has higher DxPD (trend continues for all additional write patterns).
- The 5210 and enterprise-class HDD have similar DxPD when write traffic is 90% 128K sequential; the 5210 has higher DxPD for 100% 128K sequential writes.
- The 5200 ECO DxPD exceeds both HDDs for any write pattern shown.

Using DxPD to Help Find the Right Drive – SSD or HDD

There are three key aspects of this information that will help align specific enterprise workloads to storage devices.

DWPD is Essential in Understanding Endurance in the Age of QLC SSDs

DWPD and DRWPD are ratios of how much of a drive's capacity can be written (SSD) or read/written (HDD) per day based on the drive's usable capacity, so it can be misleading to choose based solely on DWPD values. To best understand how much endurance one is getting, DxPD should be converted to gigabytes written (or read/written for

HDDs) per day. Gigabytes is simply a function of DWPD: If a 7.68TB SSD has a DWPD rating of 0.5 DWPD, it means that 50% of the drive can be written per day (3.84TB).

This is especially relevant when evaluating a QLC SSD since QLC technology packs 33% more bits into every cell and QLC SSDs are typically only available in higher capacities (which inherently have lower DWPD ratings).

For example, a 960GB TLC SSD with a 1 DWPD rating delivers DWPD similar to a 1.92TB QLC SSD that has a 0.5 DWPD rating for a given workload. While the QLC SSD DWPD specification *appears* lower, the amount of writes per day is similar.

DWPD Needs Are Lower for Many High-Growth Workloads

There is a strong trend in high-growth applications showing that they read far more data than they write. Many industry analyst firms have indicated very high growth rates for read-centric enterprise workloads ranging from artificial intelligence (AI), machine learning (ML), big data analytics, low-ingest Ceph block/object storage, some NoSQL workloads (profile caching, read latest, etc.) and deep learning and business intelligence.

These read-centric applications can be a very good fit for SSDs (which incur very little wear when written). However, if you are replacing drives based on warranted endurance values, you may want to be very selective when using HDDs for these same applications since many HDDs incur wear when read and written. While some HDDs have no workload limit rating, many high-capacity enterprise HDDs have this rating, so care must be taken with HDD selection.

Understanding Endurance

Many read-intensive workloads still have write I/O as part of their read/write ratio, even if they’re classified as being heavily read-intensive (for example, ingesting data into HDFS for subsequent analytics). SSDs can meet and in some cases, exceed those needs, even though the SSDs are classified as read-intensive drives. Read-intensive simply means that a greater proportion of the target workload involves reading data, not writing data.

Many read-intensive applications are still deployed on HDDs instead of SSDs, and this may be due to initial endurance concerns. As we’ve seen, the endurance of the 5200 ECO and the 5210 ION SSDs exceeds that of both archive-class and enterprise-class HDDs under common use cases (depending on the workload I/O profile).

Conclusion

SSDs and HDDs (with workload limit ratings) incur wear differently. SSDs wear when written, and their wear tolerance is expressed as DWPD. HDDs are different. HDDs with a workload limit rating incur wear when read or written. In this paper we express their wear tolerance as drive reads or writes per day (DRWPD).

Their capacities differ as well. Converting DWPD and DRWPD into GB/day values helps normalize the effects of capacity and makes it easier to compare their wear differences.

We noted earlier that when any drive is deployed in an IT environment — HDD or SSD — the choice depends on several factors that may vary widely among organizations. Your company may prioritize time, acquisition cost, TCO or other factor.

Understanding your workload by using analysis tools built into many operating systems can also help. These tools can show you what I/Os are being sent to storage and provide a deeper understanding of how your applications are using storage and whether the application is more write-intensive or read-intensive.

If you make drive replacement decisions based on rated endurance, we showed that for many workloads, SSD DWPD ratings are equal to or exceed the DRWPD of some capacity-focused, enterprise HDDs — making SSDs a great fit for emerging and traditional read-focused enterprise workloads.

micron.com