

Matching SSD Endurance to Common Enterprise Applications

Selecting the correct SSD for your enterprise application can seem like a complicated process. The first things to consider are form factor, interface, and capacity. After those main choices are made, the final choice impacts both your acquisition and operational costs: SSD endurance. This paper explains why SSD endurance is important and how Western Digital specifies it. There are some rules-of-thumb for selecting the correct SSD endurance level for your application.

SSD Endurance in a Nutshell

The underlying physics regarding SSD endurance are complicated, but the results are simple: SSDs are built using NAND flash, and the constituent NAND flash cells can be erased and rewritten only a limited number of times before they become unreliable. The industry uses the term program/erase cycles to describe events that take place within NAND flash to rewrite data within a NAND cell or group of cells. As the industry transitions from Multi Level Cell (MLC) to Triple Level Cell (TLC) and Quad Level Cell (QLC) SSDs, which store 3 and 4 bits per cell respectively, the number of times the NAND cell can be rewritten (P/E cycles) decreases dramatically.

SSD write endurance, often referred to simply as "endurance," is the total amount of application (operating system and file system) data that an individual SSD is guaranteed to be able to write under warranty. Fundamentally, the endurance of any manufacturer's SSD is a function of the raw NAND flash cell P/E cycle capability combined with the technology that the manufacturer puts around it. Western Digital is no different.

Western Digital improves NAND flash longevity in two ways: error correction and overprovisioning. Advanced error correction techniques such as our CellCare™ NAND management technology or Guardian Technology™ platform can extract good data from even marginal flash cells and can dramatically extend the NAND cell's usable write lifetime. Overprovisioning adds additional flash capacity to the SSD, which is not visible to the user but extends endurance by allowing for more efficient data management and distributing total number of writes across a larger population of NAND flash.

The error correction technology and the amount of overprovisioning that Western Digital incorporates into an SSD can vary by product line. Users can select between different endurance levels to optimize the amount of data written and the acquisition costs for their specific applications. Note that endurance levels between product lines vary, and even within a single product there are often multiple endurance levels from which to choose, such as the Ultrastar® DC SS300 SAS SSDs, which are available 10 DW/D, 3 DW/D and 1 DW/D endurance levels.

Endurance on the Datasheet

The endurance of an enterprise SSD should be provided on its datasheet in one of two forms: either "n DW/D" (n drive writes per day) or "m TBW" (max terabytes written). While the "terabytes-written" form is relatively straightforward because it represents the total amount of data that can be written to the SSD over its warranted lifetime, the "drive-writes-per-day" form can be slightly confusing.

The term "drive writes per day" explains how many times the total usable capacity of an SSD can be overwritten within a 24-hour period over the warranted lifetime. This DW/D value can vary between 0.1, "read-optimized," all the way up to 10, "write-intensive," or higher. If, and only if, two drives' usable capacities and their warranties are identical, can you then directly compare their DW/D figures to determine which has a greater write lifetime. Conversely, comparing DW/D values between drives of different sizes or warranty periods is not possible without conversion to TBW.

SSD endurance values on datasheets should also indicate the workload for which this endurance is applicable, either small block (i.e., 4K or 8K) random or large block sequential. When both are specified, the sequential endurance value is often much higher than the random value, simply because sequential workloads are not subject to as much overhead (also known as write amplification) as random workloads. A good example of this is the Ultrastar DC SA210, an enterprise-class SSD that is optimized for content distribution workloads and provides 0.7 DW/D for sequential use or 0.1 for a more random workload. The Joint Electron Device Engineering Council (JEDEC) offers standardized workloads for endurance testing that are detailed in JESD219 as either Client or Enterprise workloads, the latter of which is a fully random data pattern touching every logical block address (LBA) intended to be run 24 hours a day for the warranty period. It is essential to verify whether or not the data pattern used for endurance rating matches your intended workload.

Converting Between DW/D and TBW

As long as the capacity and warranty are available, the conversion between the two measurement units is straightforward.

$$\text{Convert from DW/D to TBW: TBW} = \text{Capacity(TB)} * \text{DW/D} * 365 * \text{Warranty(years)}$$

$$\text{Convert from TBW to DW/D: DW/D} = \text{TBW} / (365 * \text{Warranty(years)} * \text{Capacity(TB)})$$

1 DW/D Does Not Equal 1 DW/D

A common misconception when examining SSD specifications is that "1 DW/D" on one drive means the same endurance as "1 DW/D" on another drive. That's not true, even for drives of the same model but different capacities. To understand why, examine the case of a 3-DW/D, 800GB¹ Ultrastar DC SS300 SSD and a 3-DW/D, 1.6TB Ultrastar DC SS300, both of which share a 5-year limited warranty:

$$\begin{aligned} & \text{800GB Ultrastar DC SS300 @ 3 DW/D:} \\ & 0.8\text{TB} * 3 \text{ writes/day} * 365 \text{ days/year} * 5 \text{ years} = \\ & \sim 4.4\text{PB lifetime data written} \end{aligned}$$

¹ One gigabyte (GB) is equal to one billion bytes, one terabyte (TB) is equal to 1,000GB (one trillion bytes) and one petabyte (PB) is equal to 1,000TB when referring to storage capacity. Accessible capacity will vary from the stated capacity due to formatting, system software, and other factors.

1.6TB Ultrastar DC SS300 @ 3 DW/D:
1.6TB * 3 writes/day * 365 days/year * 5 years =
~8.8PB lifetime data written

This effect becomes even more pronounced as drives increase in capacity. When comparing the highest capacity model of the Ultrastar DC SS200 at 1 DW/D with the highest capacity Ultrastar DC SS200 available in a 3-DW/D specification, we see that even though the DW/D value of one is three times the value of the other, the actual amount of lifetime data that can be written to the 3.2TB model is only about 25% more:

Ultrastar DC SS200 7.68TB @ 1 DW/D:
7.68TB/day * 365 days/year * 5 years =
~14PB lifetime data written

Ultrastar DC SS200 3.2TB @ 3 DW/D:
3.2TB * 3 writes/day * 365 days/year * 5 years =
~17.5PB lifetime data written

Choosing the Wrong Endurance

Over-specifying the endurance requirements for an SSD is a relatively benign error. At worst, it increases the initial acquisition cost but should not affect ongoing costs. In fact, in some cases the higher-endurance SSD can provide even higher write performance if your application can take advantage of it.

Conversely, under-specifying the endurance requirements can have a serious impact, both on reliability and ongoing costs. Users are encouraged to monitor a drive's endurance usage on an ongoing

basis. All Western Digital enterprise SSDs provide lifetime wear statistics to help identify when a drive is close to using up its rated endurance. As the total amount of written data goes beyond the warranty endurance, the possibility of data loss and drive failure increases. In this case, failing drives will need to be replaced, and corrupted or lost data may occur.

Selecting the Right Endurance SSD

Given the dangers of under-specifying the endurance, carefully selecting an SSD of the proper endurance level maximizes reliability and minimizes costs.

For workloads in which the amount of written data is well known, either by design or by measurements of existing systems, the choice of endurance levels is straightforward: Determine the average amount of data written per day, multiply it by the number of days a server is in service (i.e., three or five years), and then use that number as a lower bound for your endurance selection. This average number is a lower bound because when using measurements of today's workloads and extrapolating, it's a prudent practice to add headroom for expected growth.

When workload measurements aren't readily available, there are some rules of thumb that can be used instead. Table 1 contains a list of use cases and a generalized range of DW/D, adapted from our paper [Top Considerations for Enterprise SSDs](#). As these rules are not hard and fast, being the result of conversations with customers and our product teams, please consult with your sales representative to ensure that they make sense for your specific use case.

Use Case	Description	Approx. DW/D	SSD Models
Boot Drive	Server boot drive. Updated only periodically. Logs and all permanent data stored elsewhere.	0.1 ~ 1.0 DW/D random	Ultrastar DC SA210, @ 0.1 DW/D
Content Distribution	Accelerating CDN front ends. Media migrated from HDD to SSD on an hourly basis depending upon popularity.	0.5 ~ 2.0 DW/D sequential	Ultrastar DC SA210, @ 0.7 DW/D
Surveillance	Streaming writes from multiple cameras. These writes operate continually, and the entire drive is overwritten on a periodic basis.	Calculate directly: Cams * BW sequential	Ultrastar DC SS200, @ 1/3 DW/D
Virtualization and Containers	Tier-0 storage for containers and VMs in a hyperconverged system. SSDs provide all local storage for the cluster.	1.0 ~ 3.0 random	Ultrastar DC SS200, @ 3 DW/D
OLTP Database	Data intensive workloads. Frequent updates to database logs and data files, often thousands of times per second.	3.0+ random	Ultrastar DC SS200, @ 3 DW/D Ultrastar DC SS300, @ 3/10 DW/D Ultrastar DC SN200, @ 3 DW/D
High Performance Caching	Accelerate local hard drives. Some of the highest write workloads possible (since even reads can turn into cache writes).	3.0+ random	Ultrastar DC SN200, @ 3 DW/D Ultrastar DC SS300, @ 10 DW/D

Table 1. Endurance selection rules of thumb, (when measurements are unavailable)



Conclusion

Selecting the correct SSD for your applications requires choosing the appropriate endurance. Data sheets typically present endurance levels in one of two ways, DW/D or TBW, and, under specific workloads, they're presented as random or sequential. By using information presented in data sheets along with either real-world measurements or general rules-of-thumb you can select an SSD with the appropriate endurance level for your specific needs. Moreover, the proper endurance level will maximize your SSD's lifetime and minimize your acquisition costs and ongoing expenses.

Western Digital.

5601 Great Oaks Parkway
 San Jose, CA 95119, USA
US (Toll-Free): 800.801.4618
International: 408.717.6000

www.westerndigital.com

© 2018 Western Digital Corporation or its affiliates. All rights reserved. Produced 7/18. Rev 9/18. Western Digital, the Western Digital logo, the HGST logo, CellCare, Guardian Technology and Ultrastar are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. All other marks are the property of the respective owners. References in this publication to Western Digital products, programs, or services do not imply that they will be made available in all countries. Product specifications provided are sample specifications that are subject to change and do not constitute a warranty. Please visit our website, www.wdc.com/dc-support for additional information on product specifications. Pictures shown may vary from actual products.