Use Case Impact on an SSD's Lifespan

Overview

Due to their extreme performance, low power, and reliability, client SSDs are deployed in a wide spectrum of applications including Workstations, PCs, NAS, surveillance, and even embedded systems. However, each application has a different impact on the SSD's endurance over time. SSDs are solid-state drives that are engineered with NAND, which has all the aforementioned benefits, but are prone to wear which directly affects the lifespan of the drive. Beyond meeting or exceeding the performance requirements, how do you know if the drive has enough endurance for the use case expected so the drive doesn't wear out and fail? This brief is intended to provide considerations and recommendations.

Background

NAND flash memory is a non-volatile storage that retains data without power and is the storage component in SSDs. The basic building block is the memory cell, which stores bits as electrical charges. NAND flash memory has also evolved to keep up with the trend for higher capacity SSDs and can store multiple bits per cell. As seen in Figure 1, MLC (Multi-Level Cell - 2 bits/cell), TLC (Tri-Level Cell - 3 bits/cell), and QLC (Quad-Level Cell 4 bits/cell) are the NAND types deployed in today's SSDs.

Figure 1

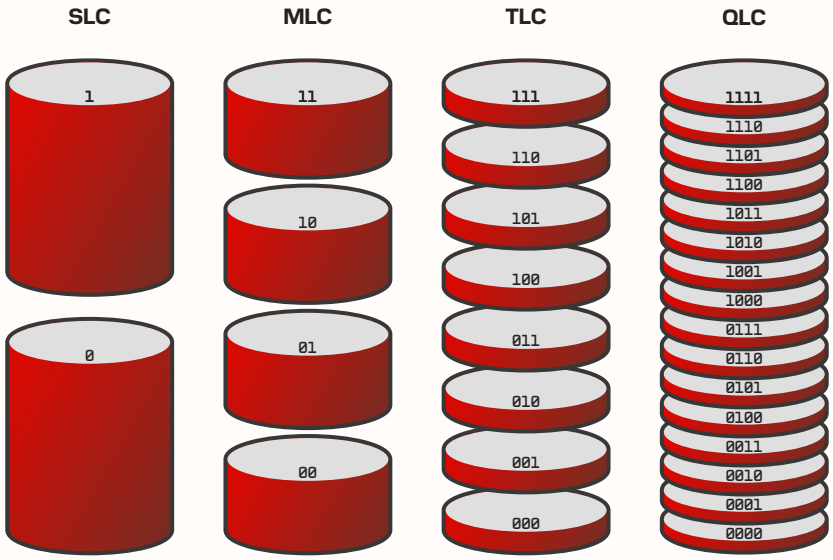

Memory cells are organized into pages (for reading/programming) and blocks (for erasing). Blocks can consist of 512 pages or more. While data can be read and written (programmed) by pages, memory cells must first be erased by blocks before writing data to the memory cell, consequently, causing writes to be slower than when reading data.

Each memory cell has a limited number of program/erase (P/E) cycles in its lifespan. This P/E limit is a factor of the type of NAND (SLC, MLC, TLC, QLC) and typically decreases as the number of bits per cell increases. To mitigate wear and tear on the memory cells, the SSD's firmware will continuously monitor each cell and rotate or replace it with spare blocks consisting of new memory cells. Each new drive comes with spare blocks for this purpose. But as the drive fills up over its lifespan or the number of spare blocks diminish due to wear, the FW's ability to prevent catastrophic failure becomes problematic.

The write amplification factor (or WAF) largely impacts wear on memory cells. WAF is the ratio of the number of actual writes to NAND divided by the host write requests. Host operating system and drivers should write in flash friendly manner such as large sequential blocks and avoid random small blocks or performing misaligned (i.e. crossing block boundaries) transfers.

Mitigating Wear

Sandisk goes through great lengths to help ensure the drive's performance and health is maintained through the drive's warranty[1] period and beyond. Based upon the expected use case for a specific SSD model, we select the SSD components to ensure an SSD meets its published specs. The following are some of the architectural ways we mitigate wear and tear on the SSD.

- We select a specific NAND type for the use case. Typically, the more bits/cell in a NAND memory cell, the less P/E cycles (PEC) a memory cell can incur. TLC has more PECs than QLC NAND.
- Our firmware (FW) continuously performs wear leveling and garbage collection. During idle time, our FW rotates blocks for even wear using spare blocks when needed. The number of spare blocks is based upon capacity and free space. High capacity drives typically have a higher percentage of spare blocks.
- We use write caches. Adding DRAM or an SLC write cache helps reduce WAF by coalescing smaller writes before data is written to user space in idle periods.
- There are a few other methods which can be deployed for specific SSDs such as integrating a special "trim" of the NAND designed for higher endurance, adding more spare blocks, or commonly used in enterprise SSDs, overprovision the SSD (i.e. reserving a % of user space).

Endurance Specification and Testing

Consumer SSD's endurance spec is typically given as a TBW[2] (terabytes written) unit which the total amount of data (expressed in Terabytes) which can be written to the SSD over it's lifespan. If this is exceeded, then the drive needs to be replaced and will commonly be switched to read-only. While an SSD's TBW spec is set by the vendor, it typically is a factor of the NAND type, expected WAF, and the SSD's capacity over a warranty[1] period. In general, the TBW spec will increase as the capacity increases. Consumer SSDs typically also include a limited warranty which states that the drive will maintain functionality until the warranty period elapses or the endurance spec is reached.[1] Common warranty[1] periods are 3 years or 5 years. When properly selected for a use case, an SSD should last much longer than the warranty[1] period.

Sandisk uses JEDEC standard for testing the SSD's endurance spec. JESD218A is the testing methodology which simulates a typical consumer use behavior allowing for idle periods for the drive to perform wear-leveling and garbage collection. This testing will use the JESD219A client workload which consists of real-world storage traces and includes a higher percentage of sequential and a smaller percentage of random workloads.

Use Case Impact on Endurance

Let's consider consumer use cases which may impact the SSD's endurance. For example, professional videographers, film studios, or commercial deployments are typically write intensive and have a much greater impact to the SSD endurance than consumers performing general productivity, web browsing, and emails. Therefore, use cases which have signficantly higher number of writes to the SSD can affect the long term lifespan of the SSD. A typical method to understand use case's impact to drive endurance is to estimate the amount of data written to the drive, every day, for a year. This gives a good gauge as to whether you're selecting the correct drive prior to purchasing.

As seen in Figure 2, professional content creator workflows can exceed 200 GB[3] written per day generating a high level of program/erase cycles on the SSD's flash memory cells. Another example of endurance stress is a 24/7 NAS system in a multiuser environment generating a wide range of data traffic including random writes across a wide range of logical addresses to the SSDs. For both cases, an SSD engineered specifically for high endurance is recommended.

Prosumers and gamers generate a large amount of writes to their SSDs during video editing, game installation, patching, and play where high performance also must be sustained.

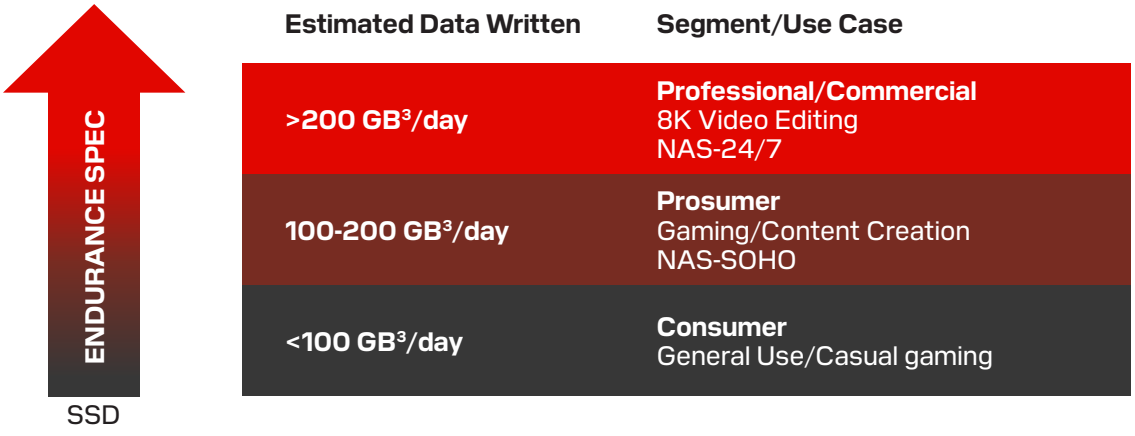

Figure 2

Sandisk SSDs

Sandisk has one of the largest client SSD portfolios to match major segments including professional, NAS, gaming, prosumer, and consumer. Sandisk drives are engineered to meet or exceed endurance requirements for those segment's use cases.

Our SSD portfolio includes WD_BLACK™ for gamers, WD Blue® for Creators, WD Green® for Everyday consumers, and WD Red® for NAS. As seen in Figure 3, a WD_BLACK SN850X SSD deployed in a professional workflow can be an ideal choice as endurance at every capacity can meet or exceed the use case's requirements for the 5 year limited warranty[1] included.

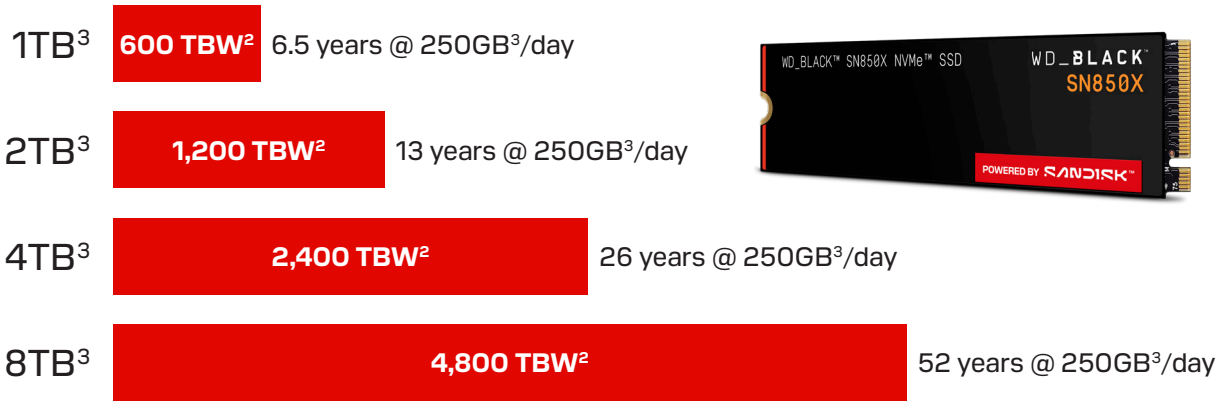

Figure 3: WD_BLACK™ SN850X SSD for Professionals

Recommendations

While SSDs can be engineered and tested to meet JESD219A standards for a specific warranty[1] period and consumer segment, they will have varying TBW[2] specs. A drive manufacturer selects the components that best matches the expected use case's endurance requirements and segment's price points.

Consumers should carefully assess their use case write per day requirements and budget spend prior to selecting a drive. We also recommend that consumers select a higher capacity drive if they can afford to allow additional space for the SSD to maintain performance and proper wear-leveling. Consumers should avoid filling their drive near capacity for this reason.

Professionals deploying SSDs within a 24/7 NAS or within write intensive workflow should consider SSDs with a DRAM or SLC write cache and has a maximum endurance and the highest capacity.

In conclusion, use cases can often contribute to the wear of SSD even though SSDs are designed for the long haul. Consider determining the SSD's use in the system and it's expected write impacts to storage. We recommend always buying one capacity higher than expected to allow ample space for future user data along with providing the drive's firmware with more spare blocks to keep in peak performance and health over it's lifespan.

1. Limited Warranty Period or Max Endurance (TBW) limit, whichever occurs first. See sandisk.com/support/store/warranty-services for regional specific warranty details.
2. TBW (terabytes written) values calculated using JEDEC client workload (JESD219) and vary by product capacity.
3. 1GB = 1 billion bytes and 1TB = 1 trillion bytes. Actual user capacity may be less depending on operating environment.

951 SANDISK DRIVE
MILPITAS, CA 95035, USA
WWW.SANDISK.COM