


Solid State Drive Applications Based on DWPD



SERVERS

DATA STORAGE

PCs & WORKSTATIONS

The logo for nfina, featuring a stylized white 'n' symbol followed by the lowercase letters 'nfina' in a bold, sans-serif font.

DWPD RATING IS CRITICAL TO THE SSDs INTENDED USE



We can all agree that flash drives have revolutionized computer storage in recent years. However, there are a myriad of Solid State Drive (SSD) technologies in use today (i.e. SLC, eMLC, MLC, TLC, NVMe, Optane). The applications range from boot drives, to cache drives in hybrid storage systems, to high capacity storage, and everything in-between.

This paper highlights some differences between the technologies and how they should (and shouldn't) be used in different applications.

TECHNOLOGY

Flash NAND Based Technology

The underlying technology behind a conventional flash based storage device is known as a NAND gate. A NAND gate is an semiconductor solid state inverting building block comprised of transistors. Two NAND gates feeding back on each other is known as a flip flop. In the flip flop cell, a digital "1" or '0" can be stored. There is also 2D and 3D technology which refers to the semiconductor cell stack-up layers (manufacturing semiconductor deposition techniques) to increase the storage density. 3D NAND is the most prevalent technology being used in SSD manufacturing today.

These devices are also nonvolatile, so they can hold their storage even when power is removed. The charge is subsequently stored in the semiconductors substrate persistently. Voltage is required to charge (write) and discharge (clear) the gates. Over time, this writing and clearing degrades the ability of the flash memory to hold its value. Therefore, these SSD devices are manufactured with a rating based on the number of digital write per day (DWPD).

DENSITY PER NAND CELL

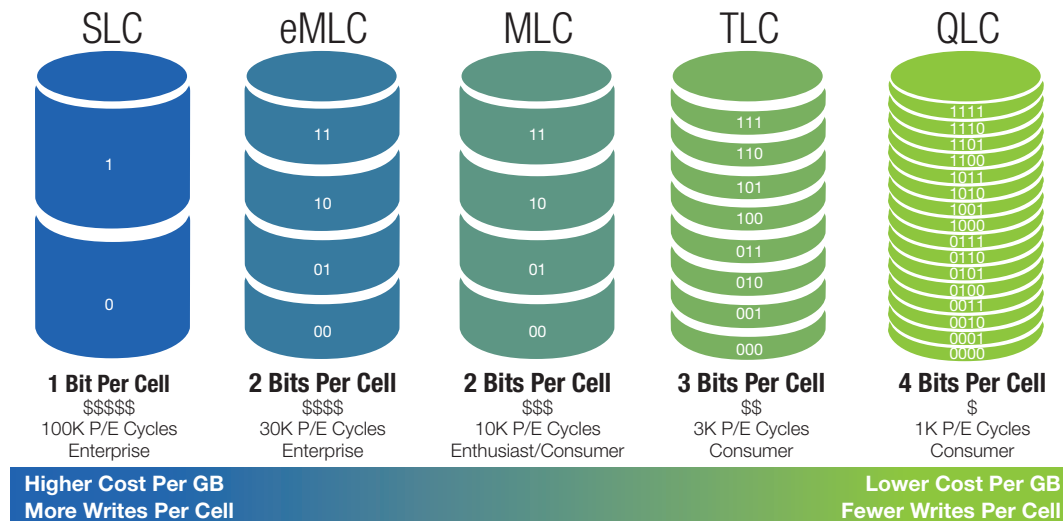


Figure 1, How to tell if a flash drive is suitable for the intended use, enterprise down to consumer.

Each type of NAND flash has a different life span meaning it can accommodate a finite number of Power/Erase (P/E) cycles before it degrades and eventually fails. NAND flash types are defined below and Figure 1 illustrates the density per NAND cell, cost, P/E cycles and intended use.

- SLC 100K P/E Cycles stores a single bit and has higher DWMP, it utilizes more silicon and is more expensive.
- eMLC 30K P/E Cycles is an enhanced (hardier) version of MLC that is in-between SLC and MLC.
- MLC 10K P/E Cycles stores two bits, has a lower DWPD rating, and is less expensive than SLC.
- TLC 3K P/E Cycles (Triple level Cell) less expensive than MLC and has 3 cells of storage.
- QLC 1K P/E Cycles (Quad level Cell) is coming soon to a drive near you to further reduce cost.

To prolong the life of the SSD, a technique is employed known as “wear leveling”. The basic technique is to oversize the storage beyond its rated value, and detect when the cell is approaching its maximum life according to its rating, and move the cells contents to one of the unused cells. Without wear leveling, highly accessed memory cells would wear out before others, resulting in premature drive failures. Figure 2 illustrates memory cell wear with and without wear leveling.

MEMORY CELL WEAR WITHOUT WEAR LEVELING



MEMORY CELL WEAR WITH WEAR LEVELING




Figure 2, Memory cell wear with and without Wear Leveling

Optane

Since Optane NVMe memory from Intel is so different, it deserves its own category. Intel calls the new Optane technology “3DXpoint” (crosspoint) that encapsulates cell addressable polysilicon cells with a metallization crosspoint matrix. It is not a transistor based design. This makes it faster and beefier than the traditional 3D NAND technology and gives it a higher DWPD yield.

It is also the fastest kid on the block sporting 500k IOPs on random read and write. The DWPD on these drives are in excess of 30 (60 for some models). The Optane memory makes for an excellent caching solution since they are an order of magnitude faster than conventional 6Gb/s SSD SATA drives. Optane is currently offered in PCIe cards, as well as in M.2 and U.2 (2.5”) formats.

Intel also has a new Optane DC (persistent Data Center) memory that upgrades memory on the PCIe bus thus making it persistent as well.

DRIVE CONNECTIVITY

SATA vs PCIe

The two prevalent connect technologies for SSD drives today are Serial ATA (SATA) and PCIe.

SATA uses the AHCI (Advanced Host Controller Interface) to connect to the drives. SATA was originally developed for spinning drives, so it is older technology and has a higher latency. Because of the lower cost and high density, SATA is alive and well.

NVMe (Non-Volatile Memory Express) SSD drives connect directly to the PCIe bus. This greatly decreases the write cycle time and results in a drastic reduction in latency (by an order of magnitude). The market share of NVMe is rapidly increasing.

SAS technology is currently being phased out for SSD drives. However, because of backwards compatibility, SAS HBA and RAID controllers can still talk to the SATA drives.

New Networking technology (RDMA, iSCSI extensions, 100Gbe, 16G – 64G Fibre Channel, etc.) can be used to greatly enhance the storage access latency of the HCI and tradition SAN storage subsystems.

Sizing DWPD to Use

For boot drives in PC's and servers, a low DWPD is acceptable. There are many vendors that market drives in the 0.3 DWPD, which is fine for this application. These drives are not acceptable for large scale data storage and/or caching high intensity write applications (NAS, SAN, File Server, backup storage, etc.).

In large scale data storage applications (i.e. NAS, SAN, Backup), a DWPD in the range of 0.65 – 1 DWPD is acceptable. If there are very write intensive applications to support, 3DWPD should be considered. Of course as the DWPD increases, so does the cost.

For the highly intensive write applications (cache), vendors recommend at least 5DWPD - 10DWPD. Of course Intel's Optane memory is perfect for cache use, and can greatly enhance any hybrid storage system's performance.

Note: One must make sure that a high bandwidth storage network (40-100 Gbe, 16G Fibre Channel, RDMA, iSCSI extensions, etc.) is deployed to exploit the difference in performance between Optane NVMe cache and tradition SSD's.

Conclusions

As shown by this paper, the DWPD rating of the SSD drive is critical to its intended use. It is imperative that users follow these general use guidelines to avoid premature failure of the drives and or the systems where they are installed.