Not All MLC SSDs are created equal

By Scott Stetzer, STEC, Inc.
1/21/2012 3:05 PM EST

With advances in NAND flash technology, solid-state drives (SSDs) have found their way into a variety of enterprise and data center applications. It is important to recognize that not all SSDs are designed the same. Some devices are made for laptops, others for industrial uses, and still others for consumer desktops owned by gamers or enthusiasts, and finally, enterprise data center uses.

The challenge is to recognize that a product designed for one use or environment is not necessarily something you should use for a different environment. For example, it is not advisable to use an SSD designed for use in laptop computers in a data center environment for mission critical applications. Why? Because laptop designs are tuned and developed for the low power requirements of a battery-driven environment with low duty cycles and will be powered up and down regularly, or shut off when not in use. On the other hand, enterprise environments expect to be powered and moving data 24 hours a day, all year, for 5 or more years, so an SSD that can survive that demanding workload is required. The opposite is equally true -- an SSD designed for the enterprise may not be the best choice in a laptop.

In this article, we focus on SSD devices intended for use in demanding enterprise environments, with a focus on the attributes, features and requirements needed to support the data center for access and reliability of user data. When selecting an SSD for these applications, IT professionals have a variety of choices available, ranging from the type of NAND flash memory technology used in the drive to how intelligent the built-in SSD controller is.

The types of NAND flash memory used in SSDs include single-level cell (SLC), multi-level cell (MLC), or enhanced multi-level cell (eMLC). While SLC flash offers high endurance and fast write access times, MLC flash stores two bits per cell to offer twice the density of SLC flash, at a more cost-effective price point. eMLC flash was designed to feature the higher densities and increased endurance that MLC enjoys, however eMLC achieves this at the expense of doubling access time, making it unsuitable for true enterprise-class applications.

SLC, MLC and eMLC Flash Memory Comparison

Metric	SLC	MLC	eMLC
$/GB	$$$	$	$+
Endurance (cycles)	100K	3-5K	30K
ECC	8b/512B	18b/512B	18b/512B

Table 1

Table 1: A quick comparison of the NAND flash memory types.

The cost-effectiveness of MLC-based SSDs (enabling more functionality in the same footprint as SLC-based SSDs) has increased their acceptance and adoption in the data center. In its raw state however, MLC flash technology has reliability and endurance challenges, especially for write-intensive and mission-critical applications. Flash memory wears out when you write to its individual cells over time, and with the array of write-intensive and mission-critical applications in the enterprise (e-commerce, online transaction processing, social networking, cloud computing, etc.), MLC-based SSDs will wear out faster than SLC-based drives, increasing a data center's total cost of ownership (TCO), unless the right tools and technologies are implemented.

By understanding how MLC NAND flash wears and identifying the different technology tools available for optimizing SSD performance and endurance, decision makers can select an MLC-based SSD that will accelerate data access in the most cost-effective and reliable manner.

Understanding Enterprise SSD Endurance

SSD endurance is defined by the measure of the usable life of the flash memory cells typically specified as the number of writes a cell can sustain. 'Writing' to a cell requires more electrical charge than 'reading' of a cell. When writing to a cell, each cell needs to be erased before it can be written to again. In either case, for every electrical charge that is passed through a NAND flash memory cell as part of a read or write operation, that cell will wear down.

In the enterprise, accelerated access to data is the primary reason SSDs are deployed, and since flash memory cells will be written to multiple times each day, endurance literally determines the reliable life of each drive. Endurance, as well as performance, reliability and availability of MLC-based SSDs are directly dependent on the design of the SSD controller (not the NAND flash memory as many suspect). The SSD controller is the brains and responds to host commands, transfers data between the host and flash media, and manages the flash media to achieve high reliability and endurance. How effectively this controller manages the flash memory will determine whether the SSD can be used in enterprise applications that require 24/7/365 uninterrupted operations under heavy read and write workloads. The real question is can an SSD manufacturer guarantee up to 30 full capacity writes per day for 5 years using MLC media to rival the endurance capability of SLC media.

A Deeper Dive into Flash Media Wear

To store data in NAND flash memory, an electrical charge is placed in the 'floating gate' portion of the NAND cell substrate which either blocks or enables electricity flow through the gate. As the NAND cell ages (or cycles), the floating gate will break down as electrons drop out of or get trapped below it. To slow the breakdown of floating gate electrons, which in turn, improves SSD endurance and reliability, enabling technology is available that slows and softens the impact that erase, write and read operations have on NAND flash memory cells. This advanced technology is described later in the article.

To prevent the NAND flash from degrading and adversely affecting SSD reliability, error correction code (ECC) technology is usually employed as a standard feature in most enterprise-class SSDs. The ECC technology enables the built-in SSD controller to detect and correct a limited number of bit errors in each block of data.

At some point, the ECC engine will be unable to correct the bit errors coming from the NAND as it wears out, so when this occurs, the SSD controller performs a read retry (to attempt to read the data again in the hope that the data is read correctly). This double layer of protection enables SSDs to have an exceptional unrecoverable bit error rate (UBER) which enables high reliability. As the NAND flash ages, the average number of read retries required will increase, and this retry will reduce the read performance, as well as the performance of the SSD over time. What is needed is an enabling technology that slows the 'wear-out' rate of the flash so ECC and retries do not need to be applied or are not significantly delayed when needed.

In reality, the larger issue in using NAND flash is the higher electrical charge used for the erase operation, and then the write operation, that primarily impacts endurance. To materially increase an SSD's operating life, more advanced techniques are required.

Techniques such as over-provisioning, throttling, compression, and de-duplication are mechanisms for delaying writes to NAND flash memory and can be effective when deployed, but actual use of these techniques does not increase the number of times to which the flash can be written. As such, these techniques are limited in the gains they can provide. Wear-leveling, for example, doesn't actually increase endurance, but instead, the flash controller spreads the writing of each data block evenly across all blocks in the SSD device to maintain consistent and even use of the NAND blocks over the life of the drive so that one location doesn't wear out faster than any other location inside of the drive.

Improving MLC SSD Endurance

To effectively improve MLC SSD endurance, advanced media management techniques need to be implemented within the SSD controller (and managed internal to the NAND flash memory). Using digital signal processing (DSP) functionality, the SSD controller can support complex management technologies such as dynamically adjusting the program and erase charges to achieve higher levels of efficiency. This enables algorithms to not only monitor how cells are being used but also predicts how reads and writes will affect performance over the long term. In this way, the algorithm can spread out cell wear to ensure consistent performance over time.

The performance, cost efficiency, and reliability of an SSD ultimately come down to the intelligence of its SSD controller. Extending the life of NAND flash memory in an SSD requires a flexible and intelligent controller architecture that can analyze and adapt dynamically to the changing characteristics of the flash cells as they age. Unmanaged MLC-based SSDs are not suitable for the continuous duty cycles and heavy workloads required for today's demanding network environments. Through the use of SSDs built with adaptive flash management algorithms and advanced signal processing techniques, IT professionals can confidently accelerate network access with near-zero downtime and advanced data protection at the lowest cost per I/O. Extending the useful life of an enterprise MLC flash-based SSD delivers a higher return on the SSD investment.

New Technology

To slow the breakdown of the floating gate electrons and to flatten MLC flash reliability and endurance concerns, technology exists today that is capable of extending the life of the NAND flash memory within an MLC SSD. At STEC, we have developed CellCare™ Technology which is based on four generations of proprietary SSD controller architecture. CellCare Technology uses a combination of write and erase management techniques, read level adjustments, write softening techniques, DSP methods for signal/bit detection, and other management technologies (as discussed earlier) to increase NAND cell life and deliver endurance benefits.

Since SSDs slow down and degrade due to the increased use of error correcting codes and retries required to overcome read issues (the more retries required, the slower the performance over the life of the drive), CellCare Technology-enabled MLC SSDs deliver fewer errors and retries reducing drive degradation while improving performance and extending its useful life (see Table 2).

Table 2 provides a quick comparison of the NAND flash memory types with MLC CellCare Technology added improves SSD endurance:

SLC, MLC and eMLC Flash Memory Comparison

Metric	SLC	MLC	eMLC	MLC w/CellCare™
$/GB	$$$	$	$+	$
Endurance (cycles)	100K	3-5K	30K	~60k
ECC	8b/512B	18b/512B	18b/512B	32b/512B

Table 2

The SSD controller transforms raw, generic flash into an enterprise-class SSD by dynamically measuring and managing flash media wear to extend SSD life without sacrificing performance. It uses advanced signal processing functions to improve flash endurance, adaptive flash management algorithms to actively manage NAND wear, and media error management and tuning functions that adjust over time and the life of the media. It also integrates other enabling technologies, as outlined in Figure 1.

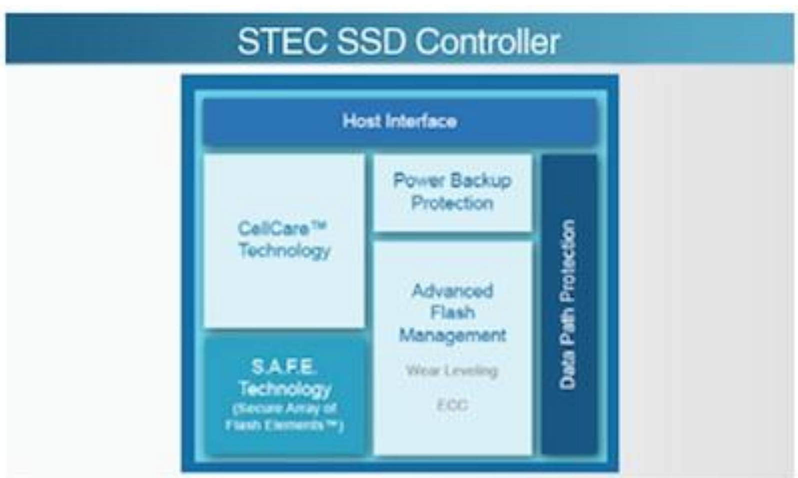

Figure 1

In addition to CellCare Technology, STEC's intelligent SSD controller includes:

Secure Array of Flash Elements™ (S.A.F.E.) Technology that improves reliability by providing mechanisms to recover from NAND flash page-, block-, die- and chip-failures while maximizing the Mean Time Between Failure (MTBF) and the Mean Time To Data Loss (MTTDL).

- PowerSafe™ Technology which provides the data persistence and instant data backup and recovery protection in the event of an unplanned power failure.

- Data path protection to protect critical data from corruption.
- Predictive read optimization that minimizes loss of performance over the useful life of the drive.
- Advanced ECC to measure and manage flash media wear.

Figure 2 shows enterprise SSD endurance based on the total drive writes per day for the NAND flash memory types based on the 32 nanometer technology node. So, the real question again is can an SSD manufacturer guarantee up to 30 full capacity writes per day for 5 years using MLC media? The answer is 'yes' as outlined in Figure 2.

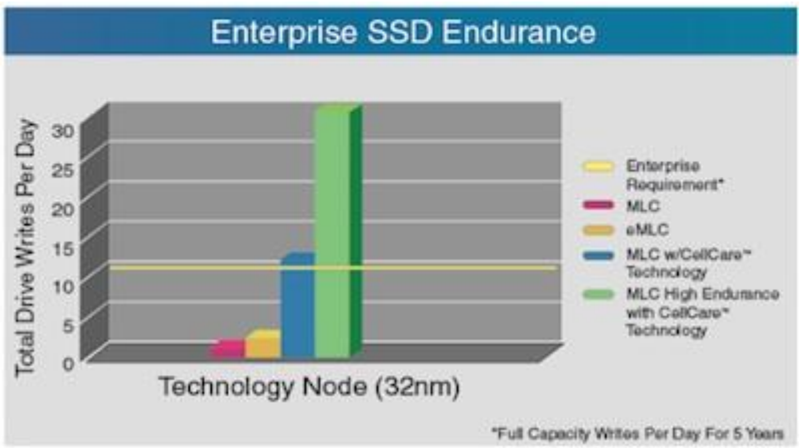

Figure 2

Summary

Selecting an MLC flash-based SSD for enterprise deployment requires more than just an evaluation of write and read performance. High-performance, enterprise-class SSDs have to be able to balance performance and endurance to provide sustained and deterministic network behavior as drives age, while providing twice the density and cost-effective price points versus SLC flash-based SSDs.

The performance, cost efficiency, and reliability of an MLC flash-based SSD ultimately comes down to the intelligence of its SSD controller. Extending the life of NAND flash requires a flexible and advanced controller architecture that can analyze and adapt dynamically to the changing characteristics of flash cells as they age. Unmanaged MLC-based SSDs are not suitable for the continuous duty cycles and heavy workloads required for today's demanding network environments. Through the use of MLC SSDs built with adaptive flash management algorithms and advanced signal processing techniques, like STEC's CellCare Technology, systems equipment manufacturers, integrators and data center managers can confidently accelerate data access with near-zero downtime and advanced data protection at the lowest cost per I/O – reducing the total cost of ownership of MLC flash-based SSDs.

About the Author

With more than twenty years of experience in storage product marketing and leadership, Scott Stetzer is Vice President of Technical Marketing for Enterprise Storage at STEC and is responsible for bringing the company's enterprise storage strategy to the industry. He has extensive storage experience in field applications engineering, software development, and technical marketing with such global companies that include Western Digital, Maxtor and Quantum. Mr. Stetzer is regarded as a futurist of storage drive technologies and has delivered keynote speeches and presentations on a variety of storage topics at key industry events that included the International Disk Forum (Japan), Storage Network World (Europe), the Data Storage Forum (China) and Flash Memory Summit (US).