

ComputerWeekly.com.com

Enterprise MLC: How flash vendors are boosting MLC write endurance

By Chris Evans

[Enterprise MLC](#) is a phrase increasingly heard in relation to [solid-state storage](#). For a long time single-level cell (SLC) was the only choice for enterprise users, while multi-level cell (MLC) was restricted to consumer devices, such as cameras, smartphones and media players.

But the old [MLC vs SLC](#) distinction is being broken down by enterprise MLC, an emerging type of [flash](#) developed by vendors working to improve MLC flash characteristics for use in enterprise applications.

In this article we will define MLC flash and look at how vendors are improving MLC flash to create enterprise MLC.

What is MLC?

In all flash memory, a [NAND](#) (negated AND) cell is used to store an electronic charge. SLC memory stores two states: 0 or 1. MLC memory stores four states and so can retain two bits of data: 00 or 01 and 10 or 11. The ability to determine four different charge states makes MLC more complex but allows for a higher density of data to be stored.

Flash devices are higher-performing than traditional hard disk drives, as latency (the time taken to start reading or writing data) is significantly reduced. Flash is expensive, but MLC is a lower-cost alternative to enterprise-class SLC.

The trade-off between SLC and MLC flash comes in performance, cost and reliability. MLC flash is cheaper to produce per bit of information stored; SLC flash is faster and more reliable, especially on write I/O.

The [I/O](#) demand on storage devices in the enterprise is significantly higher than that of the consumer market, which is why SLC was the only flash memory found in the enterprise until about three years ago. But, since 2008 vendors have started to address the weaknesses of MLC, and enterprise MLC has emerged.

Enterprise MLC and write endurance

All flash devices become less reliable over time and eventually wear out due to the way the bit state is stored in silicon. [Write endurance](#) is a measure of durability, and SLC drives are typically rated 10 times longer-lasting than MLC. For that reason the key focus for increased MLC adoption in the enterprise is in improving write endurance; there are a number of techniques vendors are using to achieve this within their products. Let's discuss those techniques.

Wear levelling. I/O is not often evenly balanced across a single drive. Write activity is usually concentrated onto specific blocks depending on the workload type. In such circumstances, MLC flash would quickly start to fail and become unusable, so wear levelling techniques have been developed that distribute I/O evenly across the entire flash drive, increasing the useful life of the device.

Two levelling techniques are in use today: Dynamic levelling writes new data to free cells, simply invalidating the old location. The other technique, static levelling, also writes data to new cells and it rebalances inactive data by periodically moving it to more active locations.

A good example of MLC flash with wear levelling comes from Violin Memory, with its [3140 Capacity Flash Memory Array](#). This product uses proprietary wear levelling technology across the whole of the array, rather than just in the individual drive. The company offers a five-year guarantee as part of its standard maintenance terms.

Overprovisioning. Another technique to increase write endurance of MLC flash is to manufacture the drive with more than the advertised capacity. In some cases, this can be as much as 20% additional space. Along with wear levelling, the increased capacity prolongs the time before the number of failed blocks makes the drive unusable. The trade-off in providing additional capacity is in the increased cost in drive manufacture.

Improved error correction. All enterprise flash devices use error correction techniques to test and correct read and write errors. As flash continues to be made to lower tolerances, the risk of bit read and write errors increases; this problem is exacerbated as manufacturers use cells to store multiple bits of data, such as in MLC flash. That's because reading multiple values from a single cell is much more inaccurate than simply reading either an "on" or "off" condition. So, to ensure reliability, more complex and efficient error correction algorithms are being produced so that MLC flash can be pushed to higher tolerances and capacities.

Write amplification and TRIM. Because existing data must be erased or moved before new data can be written to a cell, flash memory (unlike a hard disk drive) is subject to many more read/write cycles during its lifetime than are actually used. That phenomenon is called write amplification (WA). The higher the WA, the shorter the lifetime of the device since it is limited to a finite number of write cycles per cell. Optimally, a flash drive should have a WA figure of 1 (or less for MLC), but depending on workload, real-world figures can range from 3 to 20.

One way to reduce WA is via the TRIM SATA command. TRIM works in conjunction with the operating system and logically invalidates deleted data on the flash drive, removing the need to overwrite blocks as they are reused. Both Samsung and OCZ support the TRIM function, although today, operating system support is currently limited to the Windows platform.

Summary

Flash memory is finding a number of uses in the enterprise. These include:

- In boot and internal drives in servers
- As PCIe cards that emulate hard drives
- As drives in storage arrays
- In dedicated all-flash devices

Improvements in the reliability of MLC drives have enabled vendors to produce a greater range of flash-based products, including higher-capacity devices at significantly lower cost where highest performance isn't critical. Over time, we are likely to see further improvements in MLC reliability and therefore greater adoption of this technology in the enterprise.

We will also see the introduction of more hybrid technology, using SLC and MLC within the same drive and the use of flash in hybrid drives, such as the Seagate Momentus XT. The array vendors best-placed to take advantage of flash will be those offering automated tiering solutions that can operate on the sub-block level.

In future we can expect to see a continued increase in the volume of flash drives in the enterprise, with a wide mix of drive types (SLC, MLC and traditional hard disk drives) used to deliver an optimal price/performance proposition.

BIO: *Chris Evans is an independent consultant with [Langton Blue](#).*

03 Jun 2011

All Rights Reserved, [Copyright 2000 - 2024](#), TechTarget | [Read our Privacy Statement](#)