PureTruth #1: Is consumer-grade MLC flash reliable enough for the enterprise?

Welcome to a new blog series here at Pure that we're calling the PureTruth. Competitors have been known to say the darndest things, so we've decided to start this blog series to better explain common FUD items and misconceptions in the flash space. FUD, or Fear, Uncertainty, and Doubt, are little nuggets that competitors spread, sometimes maliciously and sometimes through ignorance, to try and scare customers away from competitors. In the PureTruth we'll tackle these FUD items 1-by-1, enabling you to make more informed choices about your all-flash array purchases. Enjoy!

The FUD: Consumer-grade MLC Flash isn't reliable enough for the Enterprise

The Perpetrators: Competitors who use eMLC or SLC flash in their arrays, typically found in legacy disk arrays which were retrofit with flash instead of being designed from scratch for flash.

The PureTruth: Pure Storage strongly believes that consumer-grade MLC flash is key for *both* achieving the price-point that will make flash disruptive in the enterprise, *and* for delivering the most reliable all-flash array.

- Consumer-grade flash is on the bleeding edge of flash cost reduction since it accounts for the majority of global NAND volume, and leveraging consumer-grade flash is key for hitting a price-point that enables broader deployment in the data center and the march towards the all-flash data center.
- Consumer-grade MLC has found wide-ranging deployments, from your MacBook Air to iPhone/Nexus/Galaxy to powering your compute experiences at Google, Facebook, and Amazon.

In fact, you could argue that the industry leaders have realized that to win in the market, they have to embrace consumer-grade technologies for scale and differentiation.

- cMLC, MLC, eMLC, and SLC flash all uses the same base flash chips, they differ in how they are programmed and the reliability technologies that are wrapped around them.
- Pure Storage has chosen to leverage the lowest-cost consumer-grade MLC flash available, and to build the reliability technologies into our core Purity Operating Environment software, enabling us to deliver these services more efficiently globally instead of relying on expensive eMLC SSDs. The core of these technologies is Pure Storage's FlashCare, which you can read more about here.
- Counter-intuitively, because of their mass-market unit volume, consumer-grade SSDs actually have a *lower* annual failure rate than eMLC and SLC SSDs, and can be the building block of a much more reliable flash array.
- If your all-flash array leverages eMLC or SLC flash, it is a canary in the coal mine that your vendor simply hasn't done the software work to properly understand and address flash, instead choosing to rely on exotic enterprise-grade SSDs, a path that will ultimately prove to be un-economic and less reliable.

The Bottom Line: All Pure Storage FlashArrays come with a base 3-year HW warranty, which can be extended up to 7 years as part of your maintenance agreement. If you happen to wear-out any of our flash during that period we'll replace it at no cost…but we simply don't believe this is possible, even if you run our device 24×7 at its max write IOPS.

The Long Answer…

What is cMLC vs. MLC vs. eMLC vs. SLC?

Flash SSDs are sold with a variety of different names and classifications to try and differentiate their performance, write endurance, reliability, and cost. What's interesting, however, is that all these drives basically contain the same flash chips – there is no different semiconductor process for making "SLC chips" vs. "MLC chips". So what's the difference?

- **SLC flash:** SLC means "single-level cell", each cell is charged and can determine a single voltage, essentially "on" or "off". Because a SLC cell only has two states, it can store one bit of data. This enables the flash cells to be used much longer into their lifetime, as only one voltage level has to be determined.
- **MLC flash (all types):** MLC means "multi-level cell" – and as the name implies, the flash is programmed to be able to determine multiple voltage levels within the cell, giving it the ability to store 2 bits of data. By using the same chips to store more data, reliability of each cell and longevity are sacrificed, meaning that MLC devices generally have many fewer write cycles (longevity) than SLC. On the flip side of the coin, since they store more data/cell, they are much less expensive.
- **MLC vs. eMLC:** A new class of eMLC or "enterprise MLC" devices are now being marketed. Again, these devices use the very same flash chips, however how they are packaged as SSDs differs. They typically contain more sophisticated controllers, have more flash over-provisioning to manage writes, and sometimes have more exotic cache layers with power-safe mechanisms.
- **cMLC:** means "consumer MLC" devices, and are SSDs which are specifically optimized for the needs of consumer devices. They are by far the highest volume part of the four types, have firmware specifically designed for PC/laptop use cases, and undergo extensive testing given their mass-market deployment. They are also the least expensive.

Why choose SLC?

At this point, almost no one is choosing SLC for solid state storage arrays, as they have become completely uneconomic. Back in 2008, when EMC first started retrofitting its legacy VMax/VNX arrays with flash, it used SLC flash, as almost no changes were made to the arrays themselves to manage the flash: the flash had to basically act like a spinning hard drive. Most other disk array vendor's first stabs at incorporating flash also involved SLC, but at this point SLC-based products have become uncompetitive in the market due to price, and are also really not necessary given the flash optimizations vendors have made.

Why choose eMLC?

As of this writing, most of the legacy disk arrays which have retrofit flash choose eMLC. eMLC has evolved to be able to provide the longevity in unoptimized systems that the vendors used to have to choose SLC for, at a more attractive price-point. Examples include EMC VMax/VNX, NetApp FAS, and HP 3Par arrays, all of which have retrofit eMLC flash into their disk-based storage arrays. You'll notice, however, that these vendors have not been able to push their designs to leverage MLC or cMLC flash, as the flash management technology necessary to make consumer MLC work with enterprise reliability requires a storage OS written from scratch for this task.

Why choose MLC or cMLC?

Simple: cost. MLC and cMLC devices have a significant cost advantage, but they achieve that with a reduction in write endurance. Most all-flash array devices which were built from the ground-up for flash have chosen MLC flash, including Pure Storage, Violin Memory, SolidFire, Whiptail, and others (interesting, EMC's XtremIO product still requires the use of eMLC flash...we'll let EMC explain the need for that design choice). For an enterprise-class array to successfully leverage MLC/cMLC, it must implement sophisticated write coalescing/avoidance technologies (sometimes including data reduction), array-wide garbage collection, and data integrity features. At Pure Storage this is the FlashCare layer within the Purity Operating Environment.

A surprising twist: Consumer MLC is *more* reliable (at least in a Pure Storage Flash Array)

Now let's get to the FUD part. Competitors like to spread FUD that an array designed with *consumer*-grade flash can't be as resilient as an array built from *enterprise*-grade flash. In our experience, the absolute opposite is true. At Pure Storage, we've achieved an unbelievably low SSD annual failure rate, much less that 0.1%/year (compare that to HDD failure rates, which can be 3-5%/year or higher). And in our testing (we continuously test a variety of cMLC, MLC, eMLC, and SLC drives), consumer-grade MLC drives are FAR more reliable than both eMLC and SLC drives. What accounts for this remarkable reliability difference? Two factors, one Pure Storage can take credit for, and one our drive vendors can take credit for:

- **High unit volume = more testing = higher reliability.** Consumer-grade MLC drives enjoy a production volume of over 10X the number of units shipped compared to eMLC or SLC drives, and it shows in their firmware. SSD have a single dominant failure mode: firmware bugs, not physical failure. The consumer drives are shipped in large bulk, meaning that they have to be more thoroughly tested, as a single support call will make selling that drive unprofitable to the manufacturer.
- **The Purity Operating Environment is *gentle* to SSDs.** Beyond better initial HW quality, a good deal of the work inside Pure Storage's FlashCare technology is an I/O virtualization layer that figures out the unique I/O fingerprint of each SSD we use, and understand how to send data

to that SSD in best way possible to maximize its lifespan, reliability, and performance. The result? We can literally take an SSD, run it past its manufacturer-designated write cycle endurance, put it in a FlashArray, and see no material difference in performance between that drive and a brand new one.

Why is using consumer MLC important?

Ultimately, we like to say that Pure Storage designed our FlashArray and Purity software not just for the consumer flash of today, but for the consumer flash of tomorrow, which will be continually cheaper, denser, and have less write endurance. Consumer use cases drive the majority of flash fab capacity, and consumer flash is always at the cutting edge of cost reduction in the flash space. At Pure Storage we're committed to delivering on our vision of Flash for All: and that means delivering an all-flash array product that from an economic standpoint can compete with and beat the economics of disk. So ultimately, price matters. Any product which was architected to use the crutch of eMLC (or SLC) flash just won't be price competitive in the long-run.

For Pure Storage the equation is simple: consumer MLC flash = lower cost + more reliability.