

STORAGE SWITZERLAND REPORT

FLASH SSD IS RELIABLE NOW



George Crump, Senior Analyst

IT planners are increasingly turning towards solid state disk (SSD) solutions for managing frequently accessed or “hot data” files. As a “Tier-0” storage tier, SSD can be reserved for the most active application data sets and consume a mere fraction of the rack and floor space consumed by spinning disk (Tiers-1,2,3). But there are some misconceptions about the reliability. If you are concerned about enterprise class SSD reliability you needn't be. Enterprise Flash SSD is reliable now.

Companies like [Texas Memory Systems](#) have been shipping two SSD options for the past few years – DRAM and Flash SSD. DRAM has long been associated with providing servers and PCs with a reliable, high performing “cache” memory for user applications. Flash memory, on the other hand, has been traditionally used to provide a fast memory resource for gaming systems, digital cameras, PDAs, mobile phones, and other consumer-based products.

This is where a lot of the assumptions about SSD reliability come from. Over the last couple of years, SSD suppliers have begun offering Flash-based SSD as a Tier-0 storage resource for critical business applications. Because of its legacy as a storage medium used primarily for consumer grade rather than industrial grade systems, some users

have reservations about employing Flash SSD in mission critical environments.

The rise of SSD as a more mainstream storage option in the enterprise can be attributed in part to the falling costs and increased capacity of Flash-based SSD in the marketplace over the last several years. While still an expensive storage medium when compared to mechanical drives, SSD can justify the added investment by helping businesses guarantee quality of service for key business applications. In addition, SSD can provide a tangible ROI based on lower infrastructure operational costs through a decreased reliance on multiple spinning (power and footprint consuming) drives and greater ease of management. All that remains is to address the concerns around reliability.

Flash memory is a type of silicon-based storage medium utilized by some SSD devices. Flash has a long track record as an excellent low cost alternative to DRAM (Dynamic Random Access Memory). Attractive attributes of Flash memory include non-volatility (doesn't require constant power to store data), it provides fast read access times, and it is much more tolerant of shock than mechanical disk drives.

The concern with Flash is that the architecture of the technology effectively places a ceiling on the total number of Program/Erase (P/E) operations, called Flash wearing, that can be performed before bit errors occur and the chip has to be retired. There are some other minor limitations to Flash, like the granularity by which data can be erased and rewritten on the memory chip (while data can read a byte or a word at a time, erase operations must be performed one block at a time); however, the chief concern is the premature wearing-out of the chip and the attendant risk to data integrity.

Early generations of Flash memory did indeed support a relatively low number of P/E cycles before large swaths of the available storage would be rendered unusable. This limit effectively made Flash viable only for applications like cell phones and cameras which require minimal P/E operations. How then can IT decision makers consider Flash as a viable, cost effective technology for managing mission critical business data?

Maturity Counts

The concerns over the reliability of Flash memory are somewhat similar to the concerns IT users had about ATA disk when it was first brought to

market earlier this decade. ATA was initially marketed as a low cost, near-line storage alternative to tape. ATA disk utilized less robust disk components than SCSI or Fibre Channel drives and as a result had much lower mean time between failures (MTBF) ratings. Today there is very wide acceptance of SATA and SAS disk technology as a complementary storage tier co-existing alongside Fibre Channel drives, either inside the same disk array cabinet or as stand-alone systems in the data center.

In much the same way, Flash based memory has evolved over the years to incorporate several high availability features which offset the inherent P/E cycle limitations of the underlying chip architecture.

MLC vs. SLC

Flash memory stores data in individual memory cells which are made of transistors. There are two common types of NAND Flash memory; Single Level Cell and Multi-Level Cell. The designation refers to how much data is stored per memory cell.

Single Level Cell (SLC) Flash memory stores one bit of data in each cell. This typically results in faster transfer speeds, lower power consumption, and higher cell endurance (greater

reliability). Basically an SLC-based Flash drive will perform faster and last longer, and as a result is the cornerstone of enterprise Flash.

MLC or Multi-Level Cell Flash memory storage gets its name because it stores three or more bits in each cell. By storing more bits per cell MLC-based Flash will achieve lower manufacturing costs, but at the sacrifice of transfer speeds, power efficiency, and lower write endurance (lower reliability). As a result MLC Flash is typically used in consumer-based applications like cameras and thumb drives.

Wear Leveling

Wear is a key issue for Flash SSD. Enterprise-grade Flash systems utilize “wear leveling”. This technique ensures that during a write operation, data will be evenly distributed across all of the available storage space. This eliminates the constant re-use of one small portion of the storage medium and significantly enhances the lifespan of the device. The density of the Flash chips can also improve wear leveling. Utilizing various wear leveling techniques and SLC chips, enterprise Flash SSDs can sustain workloads of 100K write IOPS over a 10 year life span.

Data Integrity

With respect to data integrity, Flash systems today employ advanced data management protection through a combination of ECC (error correction code) and DIF (data integrity field) management. The net benefit is that with the use of robust ECC, the lifespan of the Flash storage is greatly extended. In addition, the use of DIF within Flash memory is similar to performing full data checksums on all the information stored within the Flash cells. During a DIF, data is written to the Flash cell and then read completely back into a reserve cache area and compared against the checksum to ensure data integrity.

Predictability

Predictability may be of more value to an IT manager than any other Flash capability. Flash systems have the native intelligence to monitor and measure use of their storage cells; they can very accurately predict, well in advance, when components need to be replaced.

A common misconception is that most Flash failures are wear related when results from suppliers like Texas Memory Systems shows that they are not (at least for SLC

devices). The wear out problem is fairly easy to solve with a reserve of blocks. Also most IT professionals vastly underestimate just how much writing is needed to exhaust the reserve. The standard x Million MTBF failure rate per flash Die multiplied by many many dies make standard failure modes (chip failure, pin failure, plane failure) much more common than wear out. Texas Memory Systems gets around this reliability issue by RAIDing the flash chips on each board. This way ECC catches the expected media bit errors, the RAID catches the component failures, and wear leveling ensures that bad blocks are retired from service.

Flash based SSD systems offer storage planners a cost effective, high performance, reliable solution for introducing a Tier-0 resource into the data center. As a mature technology that has been around for over twenty years, Flash SSD has evolved to incorporate all of the same high availability features, plus additional capabilities that are present in standard mechanical disk array platforms. When leveraged as a high performance storage resource alongside or to replace mechanical disk systems, Flash SSD helps ensure quality of service for mission critical business data and enables IT planners to right-size their infrastructure to complement ongoing data center initiatives.

About Storage Switzerland

Storage Switzerland, is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>.