Optimizing Flash Memory for Content Delivery Platforms

Verivue, Inc. / 3 Carlisle Road / Westford, MA 01886 USA / +1 978 589 1188 TEL / +1 978 589 1138 FAX / www.verivue.com

Over the past several years, video and networking technologies have evolved to enable the on-demand world. Choppy, low-resolution Internet video and small cable on-demand libraries have mushroomed into an entirely different experience. Now broadband connections are delivering broadcast-quality HD video and increasingly, video on-demand content is expanding on digital TV services. There is a new reality to "television" and the expectation that content will be accessible on demand, regardless of the screen we use to watch it. TV everywhere and anywhere has become an industry objective.

Until recently, legacy servers, with their spinning disk drives, managed to keep up with these demands. But "three-screen" strategies, high-definition video, 3D-TV and broadband video services have changed all that. Bandwidth hungry video services are exploding and so too are the bandwidth requirements. The marketplace now requires new solutions to address these storage and bandwidth needs.

Storage Choices

Mechanical hard drives have made impressive improvements in storage density. However, with their spinning platters and mechanical read/write heads, they have maintained fairly flat bandwidth performance. Simply put, the physics of spindle speeds and seek times have not kept pace with the improvements in silicon technologies. For bandwidth-intensive applications, hard drives have become the anchor holding back performance.

With on-demand video advancing toward massive libraries and HD quality, bandwidth access is now a storage imperative. Timely silicon-based flash technology is now beginning to supplant traditional hard drives for video delivery applications.

Recently, the popularity of iPods, camera phones and other devices has fostered significant advancements in flash memory technology. Volume markets have driven rapid increases in flash storage capacity with steep reductions in pricing. Laptops as well as enterprise IT applications are now leveraging these benefits.

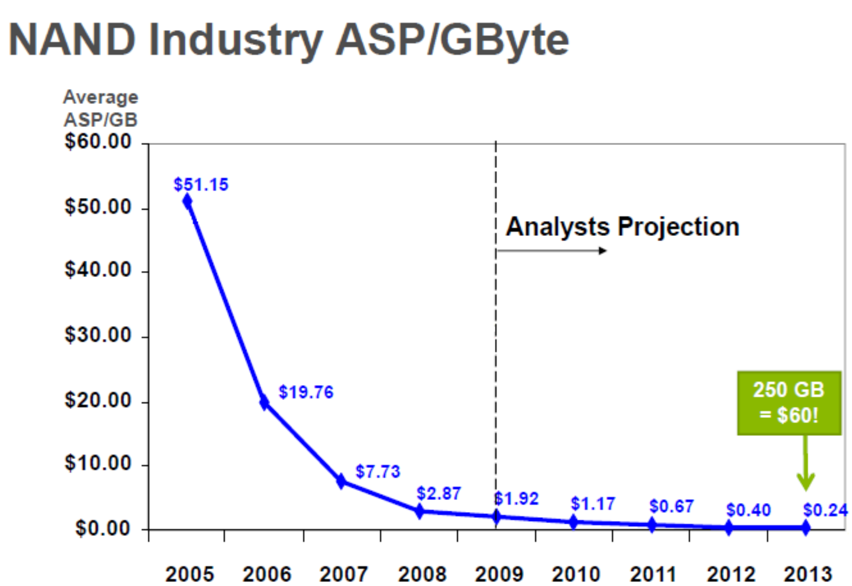

Source- SanDisk Investor Presentation, Aug 2009

Verivue, Inc. / 3 Carlisle Road / Westford, MA 01886 USA / +1 978 589 1188 TEL / +1 978 589 1138 FAX / www.verivue.com

Unlike traditional storage-intensive database applications that depend primarily on storage size, high-bandwidth applications such as video streaming benefit when the cost per bit delivered is optimized. In other words, streaming and caching applications are more concerned with the cost of bandwidth than the cost of storage. Flash memory technology delivers on this requirement.

In today's IT market, flash memory has taken a foothold within data centers through solid-state drives (SSDs). SSDs are disk drives comprised of flash memory chips (alternatively DRAM) that are packaged into an interface suitable for replacing hard-drives, such as SATA and SCSI. While still more expensive than hard drives, solid-state drives have significant benefits including access speed, reliability and lower energy costs. For many applications, these benefits outweigh the cost of reduced storage capacity. For example, companies such as EMC employ SSD storage tiers in the data center to improve I/O performance and reduce the need for dozens of inefficient 15K RPM hard drives. They find the advantages of SSD compelling to address high-bandwidth storage needs.

Flash Primer

Flash memory is non-volatile. Unlike other solid-state memory technologies such as DRAM, one of flash's greatest benefits is that it does not lose data when powered off. Although flash chips are not particularly high-speed devices, this has not been a requirement for the consumer market. While there are several variants of flash memory technologies, the majority of these storage applications use NAND flash memory. It's cheaper, denser, and provides other appropriate attributes.

Diving deeper into NAND flash memory, we find that the technology is offered in two flavors, SLC and MLC, each with different characteristics. Single-Level Cell (SLC) flash memory stores a single bit of information (0 or 1) within each memory cell. Multi-Level Cell (MLC) flash memory stores two-bits of information per cell by allowing up to four different levels of voltage to be stored (corresponding to 00, 01, 10 or 11 binary values). This allows MLC to offer higher storage densities than SLC flash parts when using the same silicon fabrication geometries. As a result, SLC flash costs at least 2x more than MLC flash for a given amount of storage.

SLC flash costs at least 2x more than MLC flash for a given amount of storage.

Due to the differing memory cell usage between MLC and SLC, the technologies exhibit different characteristics. SLC is a higher endurance chip technology, often specified in the range of 100,000 erase-cycles, while MLC targets 10,000 erase cycles. In addition, due to the higher densities of MLC, and the multiple voltage levels, MLC is more prone to errors than similar SLC flash chips.

The target volume market for SSDs is laptop drives with design goals primarily focused on cost/GB, where MLC is an attractive technology. The market is cost-sensitive and doesn't require enterprise-class reliability. Alternatively, in the enterprise market, consisting of applications such as corporate databases and email, product architects focus on reliability and write-longevity, trading off absolute cost savings as appropriate. As a result, MLC flash has been slowly adopted in the enterprise space.

Verivue, Inc. / 3 Carlisle Road / Westford, MA 01886 USA / +1 978 589 1188 TEL / +1 978 589 1138 FAX / www.verivue.com

Simply on spec, DRAM offers higher bandwidth than flash components but at a significant cost/GB. Likewise, considering basic specifications, the choice between SLC and MLC flash storage devices can be an architectural design tradeoff between storage, reliability, durability and cost.

However, with careful consideration of specific market requirements and corresponding intelligent design, solid-state memory storage can be applied to mitigate MLC limitations and leverage MLC strengths. For example, simultaneous reads and writes can be applied to parallel flash components to provide high-bandwidth capabilities that match DRAM, without its volatility. Similarly, smart engineering when applied against MLC flash can increase its reliability and durability – enabling service providers to reap MLC's clear advantages in storage density and low cost.

Maximizing Flash Durability

In content delivery applications, such as a streaming video cache, the design goals are slightly different than the general IT market. The ratio of cache reads to writes is far higher than typical IT applications.

The flash memory sub-system design becomes an issue of optimizing flash read performance while allocating a ratio of bandwidth for cache writes that satisfies this market's requirements.

Estimating that 95% of the bandwidth of the flash memory part is used for streaming, then 5% remains for content ingest. Calculating for this ratio, flash storage can be tuned to reserve 200 Mb/s of available read and about 6.5Mb/s of available write bandwidth. At this rate, a 256 Gbit flash part, which is market-available, requires 11.37 hours to be fully written. With its 10,000 erase cycle limit, this part may then be written in a sustained mode for 12.98 years before hitting its erase-cycle limit. The conclusion: MLC flash can be used to easily meet the useful lifetime of a video cache platform.

Another design technique employed to improve flash life span is wear-leveling. Wear-leveling ensures that writes to the flash chips within the storage system are spread evenly. This ensures that the 10,000 write cycles of a single MLC flash component occur across all the chips in the storage sub-system. As a result, no single chip achieves its threshold before the others, maximizing the overall storage product's expected lifetime.

Improving Flash Reliability

Flash memory exhibits characteristics called "read disturb errors" and "write disturb errors." These errors can cause bits to switch state, and therefore lose data, when adjacent memory cells are read or written. Due to the higher storage density of MLC (recall, they store 4 bits per cell), it is more likely to exhibit these errors than SLC. Many enterprise class SSDs simply choose to use SLC memory to offset these errors, however this design decision comes at a significant storage cost.

To compensate for these errors, in both SLC and MLC technology, flash memory chips and controllers may also employ error correction code (ECC) bits in their design. ECC bits allow the controller or chip to seamlessly correct errors on the fly during read operations. Flash controller design can provide varying degrees of ECC protection to compensate for MLC's characteristics -- achieving higher reliability with overall lower cost/GB.

Verivue, Inc. / 3 Carlisle Road / Westford, MA 01886 USA / +1 978 589 1188 TEL / +1 978 589 1138 FAX / www.verivue.com

Additionally, a flash memory controller can compensate for read and write disturb errors through scrubbing algorithms. As error correction bits (ECC) are used to calculate and correct for read-errors, a flash controller can pro-actively re-write the corrupted storage block once certain ECC error thresholds are achieved. This act repairs the corrupted memory cells and further reduces the chance of errors being exposed outside the storage sub-system.

VueStor™ – Verivue's Intelligent Storage Solution

VueStor™, Verivue's bladed-based storage solution, is designed for the dynamic requirements of content delivery. The patent-pending VueStor controller maximizes flash solid-state storage across all four axes: speed, durability, reliability and cost. Its architecture provides a single content store, with storage striped across storage modules and optimized flash memory chips.

Speed — Speed is maximized by storage modules with parallel access to 128 flash parts and balanced read/writes. VueStor provides significant performance advantages to all content, while supporting hot upgrades of the storage system. Importantly, as a single content store, VueStor eliminates the content replication across SSDs and HDDs that other systems employ to eliminate hot-spotting.

Durability — Each Storage Module leverages advanced wear-leveling techniques across the storage sub-system to maximize storage longevity. In addition, through purpose-built application design, write-cycles are carefully controlled, optimizing the entire storage sub-system for the erase-cycles of each flash component.

Reliability — VueStor delivers the reliability necessary for video caching and streaming. The patent-pending controller provides sophisticated 12-bit ECC and wear-leveling, guaranteed R/W performance, plus N+1, N+2 parity protection without compromising performance.

Cost — Through the proprietary hardware architecture, cost is optimized by eliminating redundant components that are required by systems that need multiple SSDs for storage and speed. Also, unlike enterprise SSDs that use expensive SLC flash memory, Verivue uses MLC technology, which provides far more storage capacity and an attractive price point. Operational benefits include ultra-low power consumption and the significant maintenance advantages of solid-state technology.

VueStor's blade-based design provides significant operational and performance benefits versus existing storage sub-systems. Supporting both 2 and 4 TB storage modules, VueStor can provide from 2 to 48 TB of storage within a single chassis. In addition, up to 12 storage modules may be striped together providing reads speeds of up to 200 Gb/s and write speeds of up to 10 Gb/s in a single, managed store.

Summary

Flash memory technology has experienced significant growth in capacity and reductions in price over the past several years. And best of all, indications are that this trend is not slowing. However, the

Verivue, Inc. / 3 Carlisle Road / Westford, MA 01886 USA / +1 978 589 1188 TEL / +1 978 589 1138 FAX / www.verivue.com

application of flash technology requires forethought and careful design optimizations for each application market's requirements.

The Verivue MDX 9000 platform and its VueStor storage sub-system was engineered specifically for today's streaming media applications. The MDX 9000 innovation brings together integrated server, switching and storage sub-systems into a single chassis in order to provide a consolidated platform for scalable content delivery. With VueStor's purpose-built design, low-cost flash technology provides unprecedented performance, reliability and longevity for video delivery applications. The result is DRAM-like performance, hard-drive capacity, solid state reliability, and the cost advantages of MLC flash's high-volume market. It's the pragmatic solution to service providers' dynamic content requirements.

About Verivue

Founded in 2006, Verivue develops and markets high-performance IP platforms that enable network operators to distribute, deliver and manage enormous amounts of IP-based media traffic across a growing and diverse number of end devices; offering solutions for Internet video, cable video and IPTV applications.

Verivue, Inc. / 3 Carlisle Road / Westford, MA 01886 USA / +1 978 589 1188 TEL / +1 978 589 1138 FAX / www.verivue.com