



EMC Symmetrix DMX to Offer Flash Drives for Ultra-High Performance

Analyst: Michael Fisch

Management Summary

Like a new island forming from the eruption of an undersea volcano, a new ultra-fast tier of storage is emerging in enterprise environments. The technology driving this sea change is *flash memory*. You may have heard of it, since flash is commonly used in consumer devices like USB drives and digital cameras, along with mobile PCs. Now EMC has caused an explosion by announcing flash drives in its high-end storage platform *Symmetrix DMX-4*. It is the first and so far only high-end storage vendor to do so.

EMC's flash drives have a special parallel architecture and intelligent media management that enable them to be robust and have exceptionally fast throughput and application response time. They deliver:

- **30 times more IOPS¹ than 15K RPM Fibre Channel drives** – i.e., one flash drive is equivalent to the performance of 30 disk drives
- **10 times faster response time for a cache-miss read**, which is on par with a cache hit
- **98% less power consumed per IOPS**

Flash drives satisfy the need for speed in applications requiring the highest available performance, such as financial trading and specialized government applications. The most relevant cost metric for these workloads today is *cost per IOPS* and flash drives can deliver the best results on that score.

The flash drives have the form factor and interface of a Fibre Channel hard drive and come in 73 and 146 GB capacities. **By integrating them into the Symmetrix DMX, EMC effectively creates a new class of ultra-high performance “Tier 0” storage that is fully integrated with the platform’s advanced data protection and management features, such as SRDF, TimeFinder, QoS management, virtual LUNs, cache partitioning, and virtual provisioning.**

In addition, EMC announced support for *virtual (or thin) provisioning* in the Symmetrix DMX-4 and DMX-3. This feature presents an application with more capacity than is physically allocated in the storage system, and allocates capacity incrementally as data is written. The benefit is lower consumption of power, cooling, and floor space, easier provisioning, and higher capacity utilization. Virtual provisioning is a foundational feature to ease and speed provisioning across all storage tiers and within DMX. Read on for details about the new capabilities of the Symmetrix DMX.

IN THIS ISSUE

➤ The Need for Speed	2
➤ Flash Memory for Enterprise Storage...	2
➤ Flash Drives for Symmetrix DMX	3
➤ Virtual Provisioning.....	4
➤ Conclusion	4

¹ Input/Outputs Per Second.

The Need for Speed

A discussion about flash memory for the data center has to begin with the need for exceptionally fast storage performance. Some applications, such as online transaction processing for financial services, require a very high throughput and very fast response time to complete their work in a timely fashion. These are situations where time really is money and, accordingly, enterprises are willing to pay more for it. Today, they currently use one of several techniques.

- **Stripe data across many disk drives** – By aggregating the performance of numerous disk spindles in a large array, it is possible to achieve high throughput (MB/s), though response time is still limited by physical drive characteristics. Striping data also requires laying out and managing a special array configuration. While performance objectives may be met by this method, it is not without significant penalties of cost and will not improve application response time. Typically, striped drives use only a small portion of their total capacity. In addition, this approach tends to consume the most power and cooling per usable GB or per usable I/O operation.
- **Lock files into cache** – This takes advantage of the fast electronic memory (SDRAM) in the front-end cache of a storage array. Cache is expensive storage and there are limits on the maximum amount of total cache in an array. Locking in files effectively limits cache resources available to other workloads. While delivering on IOPS and response time requirements, this method is severely limited in scale, if only because the maximum amount of SDRAM in an array is restricted, and SDRAM is prohibitively expensive for most applications.
- **Employ solid-state disk** – This is a stand-alone storage device that uses SDRAM instead of disk drives, similar to cache. Solid-state cache systems are also expensive and require managing an additional, segregated storage pool, one that is managed separately and is not able to leverage any of the advanced functionalities in the primary disk array, such as snapshots and remote replication.

All of these storage solutions have drawbacks. There is an open opportunity for a

different technology to satisfy this need for speed – if it can be less costly than SDRAM, offer fast throughput and response time, and integrate as part of a storage array. In fact, flash memory devices that can meet these criteria are beginning to emerge in the market.

Flash Memory for Enterprise Storage

Flash memory² is a form of electronic memory that has existed since the 1980s. Today, it is commonly used to store data in consumer products, such as USB drives or memory sticks, digital cameras, mobile phones, and digital music players. In fact, you probably use flash memory frequently whenever snapping a photo or making a call. Some mobile PCs have recently initiated its use in place of conventional drives for improved performance and battery life.

Flash memory is a semiconductor device with a significantly greater storage density than its electronic sibling SDRAM. Like a disk drive, flash is persistent, so it retains data even when power is off. It has fast performance. Depending on the implementation and the metric considered, high-end flash can be at least an order of magnitude faster than a disk drive. It also has low power consumption. At a time when green computing and reduced energy consumption receive top billing for data center requirements, flash has clear advantages.

The price per unit capacity of flash today lands somewhere between SDRAM and disk drives. In the future, the price gap between flash and disk drives will likely shrink, due to increased supply and continuous improvements in semiconductor process and design.

With these advantages, why has flash memory not been deployed in enterprise storage environments until now? There are also several challenges that an enterprise-class flash storage device has to overcome to be viable.

- To aggregate a large number of flash memory chips and achieve fast random access, a flash drive must have a sophisticated parallel architecture.
- Internal buffering – NAND flash is slower than DDR SDRAM. Often in Solid State Devices (SSDs) writes are buffered to internal SDRAM to improve write response times

² More specifically, NAND flash memory.

(which are notoriously slow on non-enterprise-class SSDs).

- To achieve device longevity, a drive must incorporate intelligent controller logic to prevent, detect, and correct bit errors, which can occur in flash memory after a relatively high number of writes. (This same issue exists in hard drives and is effectively managed.)
- Ideally, a flash memory device should integrate seamlessly into an enterprise storage array, so it does not have to be managed as a separate device and can leverage broadly used array-based software functionality.

Flash Drives for Symmetrix DMX

With its announcement, EMC has addressed these challenges and has announced support for an ultra-fast flash drive in the *Symmetrix DMX-4* platform, with availability planned for March 2008. It is the first and so far only high-end storage vendor to offer this. EMC's flash drive has a parallel design and intelligent controller logic to ensure speed and reliability. Its form factor and interface are the same as a 3.5 inch Fibre Channel drives used in Symmetrix DMX, so it plugs into a standard drive slot. EMC's flash drive packs a big performance punch in a small package.

EMC worked closely with its flash drive supplier to engineer the device from the ground up to withstand the intense workloads of high-end enterprise storage applications. Leveraging single-layer cell (SLC) flash technology and advanced controllers, they delivered on all of the enterprise requirements for ultra-fast performance, data integrity, and high reliability. Fundamental to converting the flash technology into a viable solution was the crucial step of optimizing the code in its *Enginuity* operating system to integrate the device into the Symmetrix DMX storage platform. As a result, the DMX treats and manages the flash drive like any other. It can apply the DMX's advanced software features, such as *SRDF* and *Time-Finder* for remote and local replication, QoS management for prioritizing workloads, virtual LUNs, cache partitioning, and virtual provisioning. (*More on this to follow.*)

Moreover, EMC's flash drive enables another storage tier within the DMX platform. Previously, tiered storage meant high-performance Fibre Channel drives for Tier 1, high-

capacity SATA and Fibre Channel drives for Tier 2, and software features for protecting and optimizing the tiers. **With its flash drives, the DMX now has a supercharged Tier 0 for applications that need this level of performance.**

Key Metric: Cost per IOPS

Even if the acquisition price of a flash drive is equal to its performance equivalent of 30 or more Fibre Channel drives, this would still make economic sense for applications that need high performance. The dramatic savings on power, cooling, and floor space would lower the total cost of ownership. A flash drive also avoids the management overhead of a separate storage device, as in the case of solid-state cache systems or a special hard disk layout.

The key metric to consider for high-performance storage requirements is *cost per number of IOPS*, as opposed to *cost per unit capacity*, i.e., \$/TB, a common metric. While today many buyers with ultra-performance requirements (such as those in financial services and select government agencies) buy on \$/IOPS, the majority of traditional enterprise buyers think in terms of *cost per TB* and set out to minimize this sum. **Given the new availability of EMC's flash technology in mainstream storage, when performance matters, cost per IOPS is the relevant metric you need to consider.** It reflects the cost to do the work that the business requires. By way of analogy, miles per gallon might be relevant metric for someone who wants to drive cross-country, but if the goal is to win a *Formula One* auto race, acceleration speed and cornering are more important. The cost per IOPS metric should include upfront capacity plus operating power, cooling and floor space.

Applications for Flash Drives

Flash drives can accelerate databases, especially for hot or frequently-accessed indices, logs, and temporary tables, and thus can reduce batch-processing windows. Expected initial applications will include activities like electronic trading systems, mainframe transaction processing, OLTP, seismic processing, data modeling, and many others. IT departments usually have a sense of the applications and workloads that need to have greater performance to meet business requirements. Given the relative cost of flash SSDs, early adopters

will likely be constrained to perhaps the most performance-sensitive database tables and applications. However, as the price of flash technology falls in the future, the number of suitable applications will increase rapidly.

Virtual Provisioning

Another new feature of the Symmetrix DMX is *virtual provisioning*³, which is the ability to present an application with more capacity than is physically allocated in the storage system. As data is written, the system automatically and incrementally increases the physical allocations to keep up. The net effect is less overhead or idle capacity, because the physical pool is stretched to cover more volumes, because the need to over allocate physical storage – a common practice – just to ensure that there always will be enough, has been eliminated. EMC plans for this separately-licensed feature to be available in March 2008 across all storage tiers in the Symmetrix DMX-4 as well as DMX-3, in the *Engenuity 5773* code release. It is a foundational feature, as opposed to an add-on, meaning that it applies to all storage tiers and advanced data protection and management features in the DMX. It also simplifies the provisioning process itself, which EMC claims is significantly faster than its high-end competitors.

Virtual provisioning is a useful feature with several benefits.

- Higher capacity utilization
- Lower consumption of power, cooling, and floor space
- Simpler, faster provisioning and capacity expansion
- Defer capacity acquisition
- Reduce overall costs

Conclusion

Virtual provisioning is a great new feature for Symmetrix DMX. EMC took time to do it right, integrating it into the system as a foundational feature to ease and speed provisioning across all tiers – even if it was not first to market. The savings and easier management will please DMX customers.

Support for flash drives in the Symmetrix

DMX is revolutionary in enterprise storage. It will expand the spectrum of tiered storage in enterprise environments. Today, the industry generally thinks in terms of Fibre Channel, SCSI, or perhaps SAS drives for Tier 1 and SATA drives for Tier 2 and, perhaps, tape for Tier 3. **Flash drives – with their blazing performance, like those now offered in the Symmetrix DMX – will create a Tier 0 that complements and even displaces disk drives at the high end of the spectrum.** Applications that require very fast performance will achieve it at a lower cost per IOPS with a smaller energy and floor space footprint. The portion of data stored in Tier 0 will grow as flash memory becomes more widely adopted and flash component prices continue to fall. Moreover, look for EMC to trickle down flash drive support to its other storage platforms in the future.

These days, conventional wisdom says that EMC is more about software than hardware. In this case of flash memory, hardware still provides an opportunity for change, but the software determines what you can do with it. If you want to join this movement and enjoy the early gains of flash technology, the EMC Symmetrix DMX offers the first opportunity at the high end of enterprise storage.



³ Also known as *thin provisioning*.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Michael Fisch is Director of Storage and Networking for The Clipper Group. He brings over eleven years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at mike.fisch@clipper.com or at 781-235-0085 Ext. 211. (Please dial "211" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and "clipper.com" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "Navigating Information Technology Horizons", and "teraproductivity" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.