

EMC Symmetrix DMX-4 Ultra-Performance Tier 0 Using Flash Drives

Applied Technology

Abstract

This white paper provides an overview of Flash solid state drives as Tier 0 ultra-performance storage, available for EMC[®] Symmetrix[®] DMX-4 environments with Enginuity[™] 5773.

January 2008

Copyright © 2008 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number H4132

Table of Contents

Executive summary	4
Introduction	4
Audience	5
Flash drives	5
Usage models	6
Comparison to HDDs	6
Performance comparison	6
Flash drives for the enterprise	7
Extended use, reliability, and availability	8
Physical and power comparison	8
Configuring Flash drives in a Symmetrix DMX-4	10
Protection	10
Physical location	10
Intermixing with HDDs	10
Sparing	11
Integrated software functionality	11
Symmetrix Dynamic Cache Partitioning	11
Symmetrix Priority Controls	11
Symmetrix Virtual Provisioning	12
SRDF Remote Replication	12
Conclusion	12

Executive summary

EMC has optimized the latest release of Symmetrix® Engenuity™ version 5773 to integrate enterprise-class Flash drives directly into the Symmetrix DMX-4 storage array. Symmetrix is the first and only enterprise array with support for this emerging generation of drive technology. With this capability, EMC creates a new “Tier 0” ultra-performance storage tier that transcends the limitations previously imposed by magnetic disk drives. By combining enterprise-class Flash drives optimized with EMC® technology and advanced Symmetrix functionality—including Symmetrix Virtual Provisioning—organizations now have new tiering options previously unavailable from any vendor.

Engenuity 5773 also contains new features that provide increased storage utilization and optimization, enhanced replication capabilities, greater interoperability and security, as well as multiple ease-of-use improvements.

Flash drives provide maximum performance for latency sensitive applications. Flash drives, also referred to as solid state drives (SSD), contain no moving parts and appear as standard Fibre Channel drives to existing Symmetrix management tools, allowing administrators to manage Tier 0 without special processes or custom tools. Tier 0 Flash storage is ideally suited for applications with high transaction rates and those requiring the fastest possible retrieval and storage of data, such as currency exchange and electronic trading systems, realtime data feed processing, or mainframe transaction processing. A Symmetrix DMX-4 with Flash drives can deliver single-millisecond application response times and up to 30 times more IOPS than traditional 15,000 rpm Fibre Channel disk drives. Additionally, because there are no mechanical components, Flash drives require up to 98 percent less energy per IOPS than traditional disk drives.

Introduction

For years, the most demanding enterprise applications have been limited by the performance of magnetic disk media. Tier 1 performance in storage arrays has been unable to surpass the physical limitations of hard disk drives. With EMC’s addition of Flash drives to Symmetrix DMX-4, organizations can now take advantage of ultra-high performance optimized for the highest-level enterprise requirements. Flash drives for Tier 0 requirements deliver unprecedented performance and response times for Symmetrix DMX-4, far surpassing that of any other product on the market today.

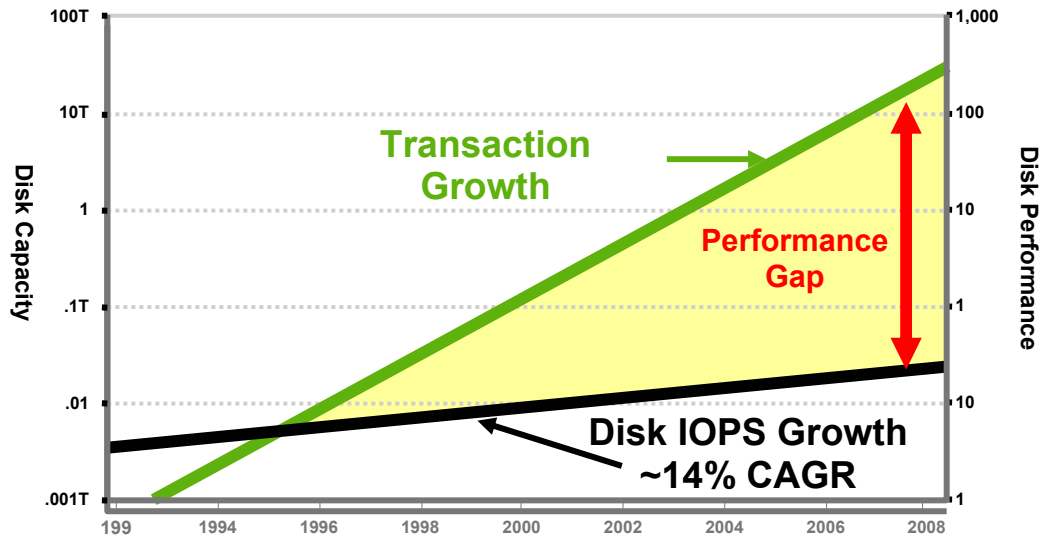


Figure 1. Disk technology lags IT performance requirements

EMC is delivering Flash drives with Symmetrix DMX-4 to provide differentiating business value to organizations whose applications require the absolute fastest response times possible. For applications such as algorithmic trading, Internet transactions, or foreign exchange, time is money—often millions of dollars

for a single customer application. Flash drives give these applications the performance they require in an enterprise-class storage array they can trust.

Because application response times of 1 ms are now possible with random read miss workloads, and lower than 1 ms response times are possible with mixed (read + write) workloads, this gives memory cache type performance characteristics to non-volatile back-end storage. In addition, the elimination of mechanical overhead and data placement latency greatly improves application performance and efficiency. Solid state technology provides up to a 10-fold improvement in response time over the fastest available HDD, and each Flash drive can deliver the IOPS of up to 30 15k rpm spindles at reasonable response times. While Tier 0 is ultra-high performance, it still utilizes standard storage management and array functionality.

This white paper explains the performance and operational characteristics of Flash drives in Symmetrix DMX-4 systems, describes the types of applications that benefit most from the use of Flash drives, and details the rules and requirements of configuring Flash drives in Symmetrix DMX-4 systems.

Audience

This white paper is intended for technology professionals who need to understand the performance capabilities and recommended uses of Flash drives in Symmetrix DMX-4 systems.

Flash drives

In a Symmetrix DMX-4 with Engenuity 5773, one Flash drive can deliver IOPS equivalent to 30 15k hard disk drives with approximately 1 ms application response time. This means that Flash memory achieves unprecedented performance and the lowest latency ever available in an enterprise-class storage array.

The enterprise-class EMC Flash drives are constructed with nonvolatile semiconductor NAND Flash memory and are packaged in a standard 3.5-inch disk drive form factor used in existing Symmetrix DMX-4 drive shelves. These drives are especially well suited for low-latency applications that require consistent, low read/write response times of less than 2 ms. As much as 30x and higher performance improvements over 15k Fibre Channel drives may be seen in mixed workload environments. The greatest improvements will be seen with higher cache read miss workloads, owing to the lack of rotational and seek latency in Flash drives.

Flash drive performance serves to increase productivity in transaction speed and volume. Write performance retains low latencies even on extremely high-write applications. This results from the unique combination of Symmetrix Engenuity, the high-performance DDR SDRAM cache in Symmetrix DMX-4 systems, and the remarkable write performance of the Flash drives themselves.

Flash drives create a new “ultra-performance” tier within a Symmetrix. This “ultra” tier benefits from the full array of advanced capabilities that Symmetrix provides, including local and remote replication, cache partitioning, and priority controls. Flash drives can even be used with the newly announced Virtual Provisioning to maximize the utilization of this new storage tier. Flash drives also integrate directly with

What is “Tier 0”?

In categorizing business and technology requirements, storage tiers are typically defined by performance, availability, functionality, and costs requirements in support of specific business applications. Of course, data can move up or down tiers as time and circumstances dictate.

Tier 1 historically demands the absolute highest performance and information availability. It also entails the highest level of local and remote replication that safeguards critical information. Examples include financial applications that sometimes process thousands of transactions per minute, or online ordering applications that are central to a business mission.

Tier 2 requirements may need only a high level of performance, availability, or functionality, which means lower cost and less critical information assets. Examples might include order fulfillment, decision support, and batch processing activities – still important, but less essential to the enterprise.

Tier 3 data includes applications like backup, archiving, and near-line storage. These typically do not need the level of performance, availability, and functionality of higher tiers.

Tier 0 is not new but for practical purposes it has been difficult to accommodate because it requires ultimate performance and lowest possible latency. Enterprise capable Flash disks change that.

The difference is that now when organizations demand even more performance for their most intensive applications, the performance increase can be gained through Flash drives supported in the DMX-4 systems. The need can be supported inside an enterprise storage system and take advantage of the management, application, and data protection capabilities that such drives offer.

storage management and Symmetrix software functionality to provide seamless service levels for critical and demanding applications.

Until now, enterprises with these high performance requirements had limited choices. They could take a costly approach of spreading workloads over dozens or hundreds of underutilized disk drives, or they could purchase separate and expensive servers and memory storage that essentially add complexity and create storage islands without capabilities typical of enterprise-class storage arrays.

Now, using Flash drives, Tier 0 applications can be closely coupled with other storage tiers within Symmetrix for consistency and efficiency, eliminating the need for time invested in manual data layout or end-of-day data transfers from separate RAM disk or specialized memory storage systems.

Usage models

Database acceleration is one example for Flash drive performance impact. Flash drive storage can be used to accelerate online transaction processing (OLTP), accelerating performance with large indices and frequently accessed database tables. Examples of OLTP applications include Oracle and DB2 databases, and SAP R/3. Flash drives can also improve performance in batch processing and shorten batch processing windows.

Flash drive performance will help any application that needs the lowest latency possible. Examples include:

- Algorithmic trading
- Currency exchange and arbitrage
- Trade optimization
- Realtime data/feed processing
- Contextual web advertising
- Other realtime transaction systems
- Data modeling

Flash drives are most beneficial with random read misses (RRM). If the RRM percentage is low, Flash drives may show less benefit since writes and sequential reads/writes already leverage Symmetrix cache to achieve the lowest possible response times. For example, if the read hit percentage is high (> 95%) as compared to read misses, such as in workloads of decision support systems (DSS) or streaming media, improvements provided by Flash drives will not likely be enough to be cost-effective.

Comparison to HDDs

Performance comparison

Fibre Channel drives and SATA disks use spinning magnetic media that store digital information. Flash drives leverage semiconductor-based block storage devices that behave as a virtual hard disk drive (HDD) via a traditional storage interface (for example, Fibre Channel, SATA).

This elimination of mechanical overhead improves performance and efficiency. Flash drives provide a 10-fold improvement in response time over the fastest available HDD. Flash drives are able to handle burst writes better than HDDs and are able to sustain lower response times even under heavy workloads.

Table 1 compares several aspects of an HDD RAID group to a Flash drive RAID group. Figure 2 illustrates performance comparisons between standard 15k rpm Fibre Channel HDDs, and Tier 0 Flash drives. All measurements were taken on RAID 5 (7+1) groups.

Table 1. HDD RAID groups vs. Flash drive RAID groups

HDD RAID group	Flash drive RAID group
Average effective application response time is approximately 5-10 ms	Consistent effective application response times of 1 ms
Performance varies based on data location and spindle contention	Little performance variance based upon data location
Fragmentation affects performance (additional seeks instead of sequential access)	Fragmentation-agnostic (random and sequential accesses are the same)
Highly dependent on workload "cache-friendliness"	Effective regardless of "cache-friendliness"




Figure 2. HDD vs. Flash drive response times

As the IOPS workload increases on the HDDs, actuator contention to service the random workload causes response times to increase to extremely high levels. Flash drives nearly eliminate this issue because they do not use mechanical components to service the request.

Flash drives for the enterprise

EMC Flash drives provide the performance and reliability expected from a Symmetrix offering. Backed by extensive research, development, testing, and code optimization, EMC has accelerated the integration of Flash technology for enterprise storage organizations.

The enterprise-class Flash drives used in Symmetrix DMX-4 differ significantly from the solid state technology used in consumer electronics, particularly in their performance and reliability characteristics.

Standard Flash-based solid state drives, such as those found in laptops, have not been optimized for writes. This means that while the read speed is much faster than a standard hard drive, the write speed is about the same. Using innovative techniques and components, the Flash drives in Symmetrix DMX-4 include a Fibre Channel interface and have been optimized for writes and longevity, giving the drive a symmetric

read/write profile. Combined with the existing high-performance cache in Symmetrix, write performance will retain low latencies even on extremely high-write applications.

The drives will buffer writes in their internal DDR SDRAM cache before destaging to Flash. The alignment between the DMX-4 tracks and the Flash buffer is leveraged by the Symmetrix write caching algorithms to minimize the number of small writes presented to the device, which serves to optimize performance and the life of the drives.

Extended use, reliability, and availability

The Flash drives also employ a technique that ensures that all cells in the Flash memory are used evenly, to minimize the risk of “wear-out” common to less advanced Flash devices. The Flash drives also safeguard data with full round-trip ECC data integrity protection and destage power backup, ensuring that the data is returned unmodified, or is identified as “suspect” much the same as does a HDD. Using advanced algorithms and block management, combined with a pool of cells that are not exposed to the host, EMC Flash drives will balance erase and rewrite operations to sustain maximum bandwidth and reliability to optimize the usable life of the drive.

With no moving parts, Flash drive architecture is not prone to mechanical failure. However, EMC still recommends configuring spare Flash drives to minimize disruption in the event of a Flash drive failure. Flash drive sparing is discussed further in the following sections.

The integrated power backup within each Flash drive is sufficient to ensure sufficient hold-up time so that any unwritten internal buffers and mapping tables are destaged to persistent Flash storage in the event of an unexpected power loss. Typically, the Symmetrix integrated SPS will be more than sufficient, but the internal protection provides an added layer of protection.

Physical and power comparison

Flash drives use less energy than rotating media, and weigh less as well. When compared to a 73 GB rotating drive, as can be seen in Table 2, the Flash drive uses 30 percent to 40 percent less energy. As the cost of these drives is reduced over time, this savings will become increasingly important.

The real energy savings, though, becomes evident when comparing the two drives based on comparable IOPS. Performance testing indicates that a single Flash drive can deliver the same IOPS as 30 of the 15k rpm drives. When this comparison is made, using Flash drives for performance can cut energy use by as much as 97.7 percent. If 120 rotating drives can be replaced by four of the Flash drives, the savings can be significant, as shown in Figure 3.

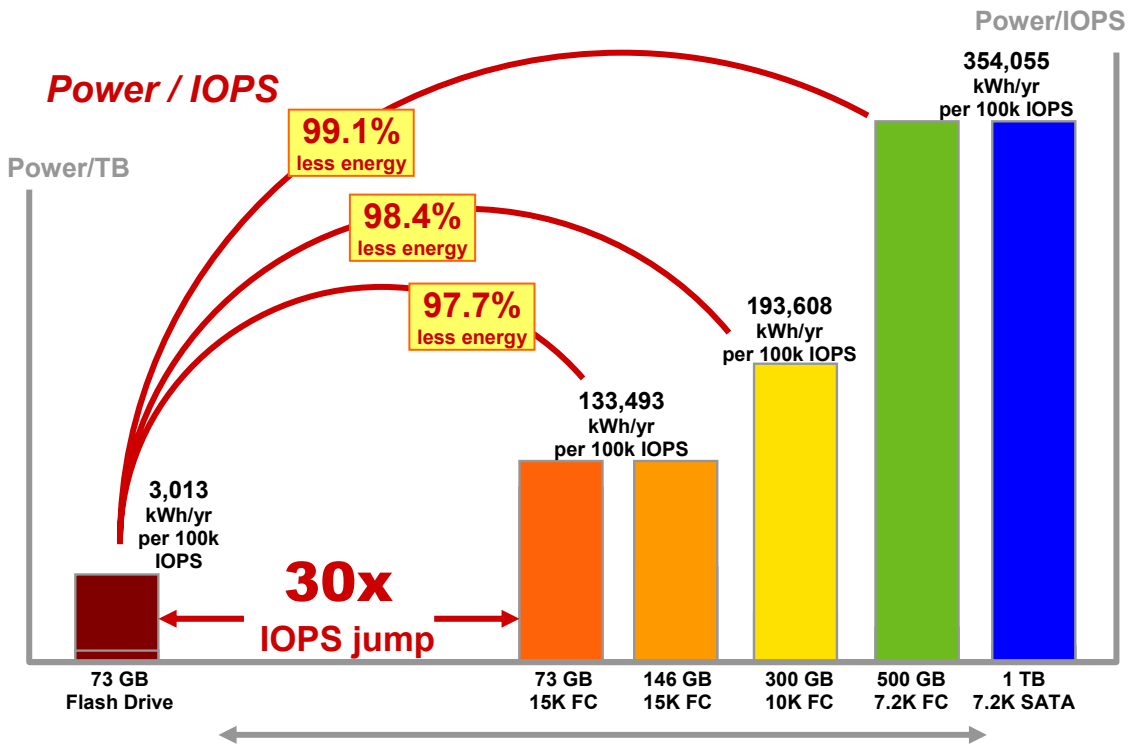


Figure 3. On a per-IOPS comparison, Flash drives require 98% less energy

And, of course, the savings in weight and floor space and even noise levels can be important as well. Table 2 compares the individual power consumption and weight of Flash drives and HDDs.¹

Table 2. Power and weight comparison of 73 GB Flash drive and 73 GB 15k rpm HDD

Drive type	73 GB Flash drive	73 GB 15k rpm HDD
AC drive power in a system	15.80 W	25.40 W
Weight	0.781 lb	1.85 lb

¹ Power numbers in the table may vary depending on exact configuration. Consult your EMC representative for specific power usage derived from the EMC Power Calculator.

Configuring Flash drives in a Symmetrix DMX-4

Flash drives are available for Symmetrix DMX-4 systems running Enginuity 5773. Existing DMX-4 and DMX-4 950 systems can be upgraded with Flash drives after upgrading to Enginuity 5773. The Flash drive is hot pluggable like a standard Fibre Channel drive. The initial offering of Flash drives in Symmetrix DMX-4 will have a native dual-ported Fibre Channel interface. The following sections describe configuration considerations when configuring Flash drives in a Symmetrix DMX-4.

Protection

Flash drives currently support RAID 5 (3+1) and RAID 5 (7+1) protection schemes.

RAID 5 is the optimal tradeoff of protection, cost, and performance for Flash drives. RAID 5 volumes must execute reads to recalculate parity on write operations. However, since Flash drive read latency is incredibly low, this penalty is negligible.

Organizations deploying Flash drives in a DMX-4 must create homogeneous RAID groups with Flash drives; all members of a RAID 5 group must be configured on Flash drives. All members of the RAID 5 group must also be configured within the same drive quadrant of the DMX-4.

Physical location

Flash drives are currently only supported in direct-connect storage bays and drive enclosures. A maximum of 32 Flash drives can be configured per quadrant.

The recommendation for a new system is to dedicate one quadrant to Flash drives. For an existing system, consider adding a new quadrant if possible or reconfiguring the system to accommodate a dedicated Flash drive quadrant.

Figure 4 illustrates the quadrants of the directly connected storage bays of a Symmetrix DMX-4.




Figure 4. Quadrants of direct-connect storage bays

In general, it is strongly recommended that Flash drive placement be balanced, configuring two Flash drives per loop.

Intermixing with HDDs

Flash drives can be mixed with Fibre Channel and SATA II disks in the same array to consolidate storage tiers, extending the tiering capabilities of Symmetrix DMX-4. In order to optimize the user benefits,

Symmetrix internal volumes (Power Vault, Symmetrix File System, and Dynamic Relocation Volumes) cannot be configured on Flash drives.

The HDDs used as vault drives may be configured with host devices, and more hard drives may be added to the loops that contain Flash drives. However, when intermixing applications on Flash drives with applications on HDDs on a single back-end loop, careful performance analysis of the specific applications involved should be conducted beforehand to prevent high-HDD activity from overcoming the benefits of Flash drives.

Sparing

One Flash drive spare must be configured within the same quadrant as the RAID group. Flash drives support both dynamic and permanent sparing using Flash drive spares. Symmetrix Engenuity will block dynamic sparing between unlike drive types to avoid potential performance issues. The standard sparing requirements for HDDs will not change if Flash drives are installed in the same quadrant.

For specific configuration rules, consult your EMC Sales Representative.

Integrated software functionality

Sophistication of the Engenuity operating environment combined with the advanced hardware technology of the Symmetrix DMX-4 is the foundation that enables effective integration of Flash drives and Tier 0 performance. Incorporating the features of Engenuity, TimeFinder[®], and the Virtual LUN technology of Symmetrix Optimizer, Tier 0 can effectively become a significant business and operational advantage.

Other examples include new Symmetrix software capabilities that EMC has recently introduced that will directly integrate with and enhance the functionality of Flash drives. These include:

- Symmetrix Dynamic Cache Partitioning
- Symmetrix Priority Controls
- Symmetrix Virtual Provisioning
- SRDF Remote Replication

Symmetrix Dynamic Cache Partitioning

Symmetrix Dynamic Cache Partitioning (DCP) divides cache into multiple partitions—set per device group—that can be either static or set to automatically share unused cache to optimize and maintain more predictable application performance in a consolidated environment.

Using DCP with Flash drives delivers response times as low as 1 ms. Since Flash drives exhibit very low read latencies, the user might want to fence all volumes on Flash drives into a small cache partition. Performance tuning using DCP will optimize cache utilization with Flash drives.

Symmetrix Priority Controls

Symmetrix Priority Controls (SPC) help manage multiple application workloads by setting priority levels and providing preferential processing within the Symmetrix system for higher-tier applications that demand specific performance levels.

Using SPC with Flash drives further enhances the Symmetrix's ability to fine-tune an application to its specific storage requirements. It is not recommended to intermix more than four HDDs per drive enclosure in order to maintain the required service levels.

Symmetrix Virtual Provisioning

Symmetrix Virtual Provisioning allows Symmetrix users to present an application with more capacity than is physically allocated to it in the storage array. It is available with Enginuity 5773 and works with Symmetrix DMX-3 and DMX-4 systems. Symmetrix Virtual Provisioning lowers TCO by improving capacity utilization and optimizing tiering capabilities across all drive types, including new Flash drives. Virtual Provisioning simplifies and accelerates processes and delivers “just-in-time” capacity allocation and flexibility.

SRDF Remote Replication

SRDF[®] is the industry-leading remote replication software, with the most flexible deployment options from synchronous and asynchronous two-site replication to advanced three-site configurations like Concurrent, Cascaded, and SRDF/Star. The new Flash drives can be used with both SRDF/S and SRDF/A.

In particular, when the replication mode is synchronous (SRDF/S), the user should be aware of delays due to speed-of-light constraints that impact I/Os between the sites. For short distances the delay might be acceptable; however, for long-distance replication, write response times could be too high if the application requires low latencies. Dependent upon the write ratios, it may also be necessary to employ Flash drives on the remote side of the replication.

It is important to note that Flash drives cannot help with the write latencies inherent with remote communication links when running any type of synchronous remote replication. However, read misses still enjoy low response times.

For asynchronous replication (SRDF/A), there are no delays associated with the distance between sites, as all writes are accepted at the source site, so write latency to the host is the same as when writing to a local Flash or mechanical drive.

Conclusion

Incorporation of Flash drives into Symmetrix DMX-4 with optimized Enginuity 5773 provides a new Tier 0 performance capability for high transaction enterprise storage applications. With comprehensive qualification and testing to ensure reliability and seamless interoperability Tier 0 is supported by key Symmetrix software applications and by storage management that enables advanced management tools including Dynamic Cache Partitioning, SPC, and Virtual Provisioning.

Magnetic disk drive technology no longer defines the performance boundaries for mission-critical storage environments. The costly approach of spreading workloads over dozens or hundreds of underutilized disk drives is no longer necessary.

Symmetrix now combines the performance and power efficiency of Flash technology with traditional disk drive technology in a single array - managed with a single set of software tools - to deliver advanced functionality, ultra-performance, and expanded storage tiering options.