

EMC Symmetrix DMX-4 and Enginuity 5772 Overview

Applied Technology

Abstract

This white paper provides an overview of the new functionality introduced with the EMC[®] Symmetrix[®] DMX-4 and Enginuity[™] 5772.

October 2007

Copyright © 2007 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 H2927.3

Table of Contents

Executive summary	4
Introduction	4
Audience	5
DMX-4 hardware	5
DMX-4 950 supports mainframe connectivity	6
4 Gb/s back-end support	6
SATA II drives	7
Enginuity 5772 algorithms	8
Improved RAID 5 performance via multiple location RAID XOR calculations	8
Optimized replication synchronization for incremental changes	8
Partial sector read hit improvements	8
Accelerated configuration changes	8
Enhanced queuing structures including FICON flow control	9
Enginuity 5772 business continuity	9
128 TimeFinder/Snap sessions of the same source volume	9
Improvements in TF/Clone create and terminate times (up to 10x)	9
Synchronous response time improvement (up to 33 percent application response time improvement)	9
Avoiding copy on first write for TimeFinder/Clone target devices	10
Symmetrix Virtual LUN	10
Clone to larger target device	10
Ongoing security improvements	11
Symmetrix Management Console	12
Improve energy efficiency	13
Conclusion	14

Executive summary

The new Symmetrix® DMX-4 system is the next generation in the DMX series and extends EMC's leadership in the high-end enterprise storage market. The DMX-4 delivers immediate support for the latest generation of disk drive technologies: 4 Gb/s FC drives for high performance and SATA II for high capacity. Symmetrix DMX-4 is the first and only high-end storage system that can support both of these latest generations of disk drive technologies. The DMX-4 is based on Enginuity™ 5772, which has been optimized for maximum performance and tiered storage functional flexibility. Enginuity 5772 provides investment protection that delivers performance gains for *both* DMX-4 and DMX-3 along with information-centric security advancements via integration with RSA enVision. With the DMX-4 and Enginuity 5772 all replication and security activities are easy to manage with the Symmetrix Management Console. And you get all of these capabilities in the industry's most energy-efficient platform, making your data center as "green" as possible.

Introduction

This white paper describes new hardware and functionality available with the Symmetrix DMX-4 in addition to the 4 Gb/s back end and SATA II support software enhancements that directly meet customer functionality and performance needs.

The following items will be discussed within this paper:

- DMX-4
 - ◆ 4 Gb/s back-end fiber loops and point-to-point disk enclosures provide throughput and loop management
 - ◆ 500 GB SATA II and 750 GB SATA II or larger drives (targeted for end of 2007)
- Enginuity 5772 algorithms
 - ◆ RAID XOR calculations possible in three locations (up to 25 percent improvement)
 - ◆ Optimized replication synchronization for incremental changes (<10 percent)
 - ◆ Partial sector read hit improvements
 - ◆ Accelerated configuration changes
 - ◆ Enhanced queuing structures including FICON flow control
- Enginuity 5772 business continuity
 - ◆ 128 TimeFinder®/Snap sessions
 - ◆ Improvements in TimeFinder/Clone create and terminate times (up to 10x)
 - ◆ Synchronous SRDF® latency improvement
 - Up to 33 percent application response improvement
 - Up to 2x the distance for the same response time
 - ◆ Avoid copy on first write for TimeFinder/Clone and TimeFinder/Snap target devices
 - ◆ Symmetrix Virtual LUN
 - ◆ Clone to larger target device
- Ongoing security improvements
 - ◆ Integration with the RSA enVision compliance and security information management platform
- Symmetrix Management Console
 - ◆ Browser-based
 - ◆ Simple, efficient, responsive
- Improve energy efficiency
 - ◆ Improved algorithms enable the same performance with less hardware
 - ◆ New energy-efficient drives
 - ◆ EMC Power Calculator

Audience

This white paper is intended for technology professionals within EMC and external to EMC who need to understand the new capabilities and subsequent productivity improvements available with the Symmetrix DMX-4 series and Enginuity 5772.

DMX-4 hardware

The Symmetrix DMX-4 series features a high-performance, Direct Matrix Architecture® (DMX). This field-proven Direct Matrix Architecture provides dedicated, non-blocking interconnects between I/O directors and global memory regions. In the Direct Matrix Architecture, contention is minimized because control information and commands are transferred across a separate and dedicated message matrix. DMX architecture delivers scalable performance to meet the most demanding information access, protection, and distribution requirements.


There are two models in the DMX-4 series: the DMX-4 and the DMX-4 950. The DMX-4 shown in Figure 1 supports full connectivity to open system and mainframe hosts (ESCON and FICON). The DMX-4 950 shown in Figure 2 represents a lower entry point for DMX technology, providing open system connectivity while offering FICON connections for mainframe hosts.

The DMX-4 is the world's largest high-end storage array, allowing configurations from 96 drives to 2,400 drives in a single system. With the option of many centralized drives, Information Lifecycle Management and tiering become powerful cost-saving strategies. The DMX-4 offers economies of scale beyond all competitive arrays.



**Combinations may be limited or restricted based on configuration*

Figure 1. Symmetrix DMX-4: World's largest high-end storage array



**Combinations may be limited or restricted based on configuration*

Figure 2. Symmetrix DMX-4 950: Lower entry point for DMX-4 technology

Table 1 compares the DMX-4 and DMX-4 950.

Table 1. DMX-4 series model comparison

Features/Specification	Symmetrix DMX-4	Symmetrix DMX-4 950
Open systems/mainframe	Open systems and mainframe	Open systems and mainframe
CPU processors	1.3 GHz PowerPC	1.3 GHz PowerPC
Maximum global memory	512 GB (256 GB usable)	128 GB (64 GB usable)
Disk adapters	1 DA pair, 2 DA pairs, 3 DA pairs, 4 DA pairs	1 DA pair or 2 DA pairs (front end/back end)
Maximum back-end channel speed	4 Gb/s	4 Gb/s
Drive channels	8-64	8-16
Minimum/maximum number of disks	96-2400	32-360
Minimum/maximum raw capacity (TB)	7.04/1053.86 TB	2.35/179.9 TB
Memory directors	8	2
Usable connectivity ports	64 × Fibre Channel 32 × Gigabit Ethernet remote replication 10 × Gigabit Ethernet iSCSI 48 x FICON 64 x ESCON	16 × Fibre Channel 8 × Gigabit Ethernet remote replication 10 × Gigabit Ethernet iSCSI 8 x FICON

DMX-4 950 supports mainframe connectivity

The DMX-4 950 supports connectivity to mainframe hosts via FICON, allowing for a lower entry point to DMX technology for mainframe users. Customers have the option to obtain FC/GbE ports, enabling connection to open system servers and the use of SRDF for business continuity, as well as all FICON ports or a combination of FICON for mainframe hosts. The combinations for connectivity are:

- 2 directors
 - 4 FICON ports
- 4 directors
 - 8 FICON ports
 - 4 FICON and 8 FC ports
 - 4 FICON, 4 FC, and 2 SRDF (FC) ports
 - 4 FICON, 4 FC, and 2 SRDF (GbE) ports

4 Gb/s back-end support

The new Symmetrix DMX-4 system supports both 2 Gb/s and 4 Gb/s loop speeds. Disk adapter (DA) pairs can be configured as either 2 Gb/s or 4 Gb/s. The type of drive on the loop determines the throughput speed. The configuration program does not allow 2 Gb/s drives to be installed on a 4 Gb/s loop but 4 Gb/s drives can be on a 2 Gb/s loop, running at the lower throughput. If a mixture of drives on a loop is replaced with all 4 Gb/s capable drives, the change to 4 Gb/s throughput for the loop is an offline event.

The DMX-4 series uses a point-to-point connection to each drive in the disk enclosure. This advancement allows:

- Access, diagnosis, and isolation of individual drives during corrective action
- Individual drives to be targeted for online testing during power-up and replacement and when adding new capacity

The DMX-4 disk enclosure delivers the same level of capacity with 15 drive slots while providing additional resiliency benefits.

The point-to-point drive connection capability enables more precise back-end fault recovery. The DMX-4 disk enclosure benefits all customers deploying either of the DMX-4 models, especially those with high drive count configurations. Figure 3 shows the traditional loop passing all traffic through all drives versus direct connection to a drive, which is now available in the DMX-4 system.




Figure 3. Traditional loop vs point-to-point topology

SATA II drives

Serial Advanced Technology Attachment (SATA) drives are a low-cost, high-density option for tiered storage. SATA uses Low Voltage Differential Signaling (LVDS), which allows very high transmission rates and an encoding scheme that eliminates the need for a separate clock signal to distinguish bit and character boundaries. SATA II refers to the use of a 3 Gb/s signaling rate, double the rate introduced with the first SATA drives. This technology promises further capacity increases over time.

SATA II drives deliver 500 GB of capacity with comparable functionality to Fibre Channel drives. (750 GB or larger targeted for Symmetrix by the end of 2007.) SATA II drives and Fibre Channel drives can coexist in the same drive enclosure but because of performance differences they should not be mixed in volume protection strategies. For example, SATA II and Fibre Channel drives should not be mixed together in the same RAID 5 group.

Information Lifecycle Management (ILM) is a strategy for managing enterprise storage to meet service objectives for different sets of data at the lowest cost. ILM strategies seek to adjust to the changing value, availability, and performance requirements of data throughout its lifecycle. By investing in lower-tier storage at the right points in time, organizations are able to meet business objectives while streamlining costs. DMX-4 SATA II drives enable ILM, providing a distinct capacity-based tier of storage within this enterprise array. As larger SATA II drives eventually come to market, the value of ILM on DMX-4 with SATA will be extended even further.

Enginuity 5772 algorithms

An intelligent, multi-tasking storage array is a prerequisite for today's complex, demanding, risk-intolerant IT infrastructures. Enginuity 5772 continues the evolution of self-optimizing storage intelligence and delivers the ultimate performance, availability, and data integrity. The following components of the Enginuity Storage Operating Environment (SOE) have been refined and refined again to deliver user critical performance and productivity improvements. These improved algorithms have been targeted to provide benefit for the most common or most user valuable operations and are available on both the DMX-3 and DMX-4. Note that the improvements listed within this document for the DMX-4 5772 Service Release, build upon features discussed in the white paper *Enginuity: The EMC Symmetrix Storage Operating Environment*.

Improved RAID 5 performance via multiple location RAID XOR calculations

Protecting data volumes with a RAID 5 strategy offers cost benefits but comes with a RAID parity write overhead. Recalculating and writing the appropriate RAID parity is required whenever RAID devices are updated. This is achieved by reading the existing parity and existing data, then performing an XOR operation with the updated data. With Enginuity 5772, RAID 5 XOR operations can be performed in several locations within the hardware hierarchy, depending on which location is most efficient. The Enginuity operating environment intelligently decides if XOR calculations are most beneficially performed in cache memory or at the disk level. With the choice of levels of hardware for the XOR operation and multiple parallel resources within those hardware levels, such XOR calculations for write operations ensure the highest possible performance.

Optimized replication synchronization for incremental changes

Enterprise class business continuity solutions have become increasingly affordable and are now relied upon to mitigate the risk of disruption to technology infrastructure across a broad spectrum of business applications. This broad range of applications may include workloads that have minimal updates during the replication cycle. The transfer of a small amount of incremental updates to the replication target can be less efficient than the process of transferring a large amount of incremental updates. The DMX-4 offers an efficiency enhancement for replication volumes that have a change rate of less than 10 percent. A new algorithm has been added to Enginuity that recognizes a "sparse threshold" and uses an optimized method to identify and copy updated tracks. Volumes that have greater than a 10 percent change rate continue to use the traditional copy algorithm. The availability of two separate algorithms guarantees efficiency and replication productivity, regardless of the change rate experienced by the replicated volume.

Partial sector read hit improvements

Fixed block architecture disk structures include the concept of the track and sector. Various disk level operations take place to guarantee what is read from a sector on a track is correctly transferred to the host. If read operations are a dominant component of a workload, then small improvements in these low level protocols can accumulate into a valuable performance enhancement. The DMX-4 has such protocol improvements. Enhancing the partial sector read algorithm is another improvement in a list of improvements that are specifically aimed at increasing user performance and productivity.

Accelerated configuration changes

The DMX-4 is the world's largest high-end storage array, allowing configurations from 96 drives to 2,400 drives in a single system. Manipulating in a satisfactory manner such a potentially large configuration and all the subsequent logical structures required an evolution in the configuration process. Enginuity 5772 streamlines host-initiated configuration changes to provide an up to three-fold acceleration of mapping, un-mapping, and metavolume creation operations. Other configuration activities also benefit from the streamlined process. With Enginuity 5772, changes to a configuration of even 2,400 drives can be achieved with a new order of efficiency.

Enhanced queuing structures including FICON flow control

Increased capacity available in the DMX-4 prompted code changes that were designed to optimize internal I/O queue structures, including enhancements to deliver appropriate scalability and I/O management for larger and more intricate environments. The changes allow for more complex queue consolidations that might have been experiencing workload imbalances. When such heavy but unbalanced workloads are present, Enginuity 5772 offers the potential for significant increases in system I/O rates. For example, in mainframe environments, such imbalances can occur when multiple logical control units fan into a small number of FICON ports. In these cases Enginuity 5772 can increase by up to three times the number of I/Os per second.

Enginuity 5772 business continuity

128 TimeFinder/Snap sessions of the same source volume

TimeFinder/Snap operations allow you to create space-saving local replicas, or snapshots between a source device and multiple virtual (VDEV) target devices, for testing, backups, or report generation. Snapshot copies can prove useful where multiple backups or recovery copies of a source device are taken throughout the day, with only a small percentage of data actually changing on the device. As an alternative to using multiple full-volume mirrors or clones, to back up these devices you can use TimeFinder/Snap snapshots to reduce costs and save on storage space. Through the use of device pointers to the original data, TimeFinder/Snap snapshots allow you to allocate space based on the expected amount of the changes to a device, rather than the size of the volume being backed up. By increasing the number of available snapshot sessions to 128 per source volume, the DMX-4 offers more flexibility in this backup management arena. One anticipated business model for this number of snapshot sessions is to take a snapshot at regular intervals to the 128 limit. When the full 128 snapshots are utilized the oldest snapshot is terminated and a new session created. This provides continuous data protection within the 128 snapshot cycle. Another case might be hourly snaps over the course of two days (yielding as an example, 48 TimeFinder/Snaps). The expanded snapshot limit for the DMX-4 offers opportunity for many additional productive uses for space-saving snapshot technology.

Improvements in TF/Clone create and terminate times (up to 10x)

TimeFinder/Clone operations facilitate the creation of *clone* copies of a source device on multiple target devices. The source and target devices can be either standard devices or BCV devices as long as they are all of the same size and emulation type (FBA or CKD). Once activated, the clone copy can be instantly accessed by a target's host, even before the data is fully copied to the target device. TimeFinder/Clone functionality is controlled via copy sessions, which pair the source and target devices. A copy session must first be created to define and set up the clone devices. The session is then activated, enabling the target device to be accessed by its host. When the information is no longer needed, the session can be terminated.

Enginuity running on the DMX-4 series contains optimization to the underlying control process that allows TimeFinder/Clone establish, reestablish, and terminate functions to receive up to a tenfold performance improvement. Once again, these improvements are specifically directed at increasing user productivity.

Synchronous response time improvement (up to 33 percent application response time improvement)

SRDF/Synchronous (SRDF/S) is high-performance, host-independent, realtime synchronous remote replication from one local Symmetrix to one or more remote Symmetrix systems. SRDF/S perfectly mirrors data between the local Symmetrix system and the remote Symmetrix partner. Data is written to the global memory of both systems before the application I/O is acknowledged as complete to the host, ensuring the highest possible data synchronization. Only when the write operation is complete at both locations can further host writes be accepted for the primary.

In Enginuity 5772, improvements to SRDF/S over Fibre Channel reduce latency caused by “transfer ready” protocol requirements. Previous versions of Symmetrix Enginuity required two round trips for every SRDF/S write I/O. This was the result of the source Symmetrix requirement to send a command frame and wait for the target Symmetrix array to return a transfer-ready frame before data transfer could begin. SRDF/S devices can now run in unsolicited mode. This new mode halves latency by allowing the initiator to send a command frame followed immediately by data. An entire round trip across the link is eliminated for each write I/O.

I/O size and round-trip time determine the response time improvement offered by this enhancement. When round-trip time is high (long SRDF link) or I/O size is large, this new protocol offers significant benefit. The typical customer benefit of this feature is a 33 percent reduction in response time, or at the same response time, the SRDF link distance can be doubled.

Avoiding copy on first write for TimeFinder/Clone target devices

With older implementations of TimeFinder /Clone technology, the first write to any track on the source or target device would result in the original full track of data being copied to the clone target to preserve a true point-in-time copy. Enginuity 5772 introduced a feature to reduce this overhead on the clone source volume. Asynchronous copy on first write (ACOFW) allowed a host update to proceed as normal, and managed the retrieval and placement of the original data track asynchronously to the I/O operation. With the 5772 service release, an extension to this process has been added to reduce COFW delays on writes to the clone target as well. This additional replication feature avoids COFW for the target device in the case of a full track write. Enginuity can automatically determine that in the full track write special case, there is no value in the COFW process and avoid this overhead.

Symmetrix Virtual LUN

As organizations consolidate multiple tiers of storage in the same enterprise system, data mobility becomes increasingly critical. With Enginuity 5772, Symmetrix Optimizer includes Virtual LUN technology that provides non-disruptive data mobility between storage tiers to help respond to changing business requirements. The Virtual LUN feature is based on proven Optimizer processes that now accommodate non-disruptive migrations of data between tiers of disk that vary by speed and/or capacity. Such data movement helps adjust to the changing value, availability, and performance requirements of data, placing it on lower-tier storage at the right time to streamline costs. Data can also be promoted again to higher-tier storage when performance requirements demand. With the latest release of Enginuity 5772, Virtual LUN technology allows a *device* group to be used to specify the migration devices. Also, a *disk* group can be used to define the possible disks receiving the migration. Devices can be migrated to unmapped and/or unmasked volumes on the disks specified to receive the relocated data. By stipulating unmapped or unmasked with a migration, the user will perform a swap with the unused volumes.

This online data mobility does not prevent normal business use of SRDF/S, SRDF/A, or TimeFinder, but does preclude the use of concurrent multiple SRDF or TimeFinder/Mirror copies. The initial Virtual LUN implementation addresses data movement only within the same RAID protection scheme.

Clone to larger target device

Enginuity 5772 and Solutions Enabler 6.4.2 provide the ability to perform clone operations from a source to a larger target device. This feature provides new flexibility, facilitating migration between devices with odd/even cylinder count variations. Certain factors must be considered to correctly manage the dissimilar clone sizes:

- ◆ The clone target cannot be smaller than the source.
- ◆ Concatenated metavolumes are not supported.
- ◆ Exact pairing operations must be used.
- ◆ Restore from a larger to a smaller device is not available.

Ongoing security improvements

As the nature of security risk evolves, there is an increase in the imperative for information-centric security to manage the relationship between data and the people that use this data. EMC has integrated technology from RSA, the Security Division of EMC, to deliver built-in security as part of the Symmetrix platform.

Critical security components of Enginuity 5772 include:

- The Symmetrix Service Credential, Secured by RSA, which provides strong authentication of valid service personnel identities and authorizes their actions in order to prevent unauthorized service actions on the service processor.
- EMC Certified Data Erasure, which automatically erases failed drives in conjunction with permanent sparing, before they are removed from the array in order to eliminate the risk of sensitive data loss.
- Symmetrix Access Controls, which restrict management actions to approved hosts, and user authorization, which provides role-based access control for Symmetrix administration.

A significant additional security component is the Symmetrix Audit Log, stored on the Symmetrix File System within the array, which provides a comprehensive, tamper-proof view of system events to support organizations' compliance with internal security policies and industry regulations. Specific logged events include:

- Host-based (for example, Solutions Enabler) actions, including device configuration, device masking, mapping, and local and remote replication
- Host-based access control failures when configured on host-based software
- User-based access control failures when configured on host-based software
- Host actions that are denied by Symmetrix Access Controls
- Active inline commands
- Configuration change scripts on the service processor
- Login attempts (successful or not) on the service processor
- Actions through the Symmetrix Service Credential, Secured by RSA, on the service processor, including denied actions
- Changes to Service Credential configuration data on the service processor
- iSCSI authentication failures

Logged entry contents cannot be altered, and access to the log can be restricted, for example with the Auditor role within Solutions Enabler 6.4. The audit trail functionality on the DMX-4 ensures records are complete, secure, and easy to understand.




Figure 4. Secure, easy, and tamperproof audit trail of interaction with Symmetrix DMX-4

The DMX-4 captures the following data in a tamperproof, easy to understand Symmetrix Audit Log located on the Symmetrix File System.

1. Host-initiated actions, including device configuration, mapping, masking, and local and remote replication.
2. Scripts via EMC service processor resident applications.
3. Manual service activities such as physically removing/replacing a component.

The Symmetrix Audit Log now can be integrated with the RSA enVision appliance - a compliance and security information management platform, which consolidates and analyzes security logs from across the enterprise. RSA enVision correlates events to identify hacker attacks or other security threats in real time or through forensic analysis and generates reports for compliance with regulatory requirements such as the Sarbanes-Oxley Act of 2002. The Solutions Enabler event daemon now can be configured to automatically push Symmetrix Audit Log entries to RSA enVision on a regular basis. This integration helps simplify the task of log management, saving time and resources and helping track Symmetrix events in the context of the larger environment. It also improves compliance reporting on Symmetrix activities. This integration is available with RSA enVision 3.5 and either DMX-3 or DMX-4 systems running Enginuity 5772 and Solutions Enabler 6.4. Log pushes are not expected to significantly impact Symmetrix performance.

Symmetrix Management Console

Symmetrix Management Console (SMC) helps simplify Symmetrix management tasks by providing a browser-based GUI for all major Solutions Enabler functionality. The lightweight SMC 6.0 software, which runs on Solutions Enabler 6.4, installs in minutes on a departmental Windows 2003, Linux Red Hat 4, or SUSE Linux 9 SP2 server. SMC accelerates FBA and CKD device management tasks, for example, enabling administrators to configure or replicate devices with a few clicks and keystrokes and to duplicate devices with simple templates containing device attributes such as capacity, configuration, and emulation type. SMC provides realtime monitoring and alerting on system health, director and drive status, and SRDF/A replication operations (looking at such metrics as cache utilization and cycle time). SMC also facilitates intuitive control and monitoring for tiered storage optimization software, including Symmetrix Optimizer Virtual LUN technology, Symmetrix Priority Controls, and Dynamic Cache Partitioning. SMC security features include user authentication, restricted user roles, management of Symmetrix Access Controls, and management of the Symmetrix Audit Log.

SMC management of the DMX-4 can be simple, efficient and responsive. Figure 5 displays the Symmetrix Management Console.




Figure 5. Easy management with the Symmetrix Management Console

Improve energy efficiency

Power consumption within the data center is an increasingly important consideration. Not only has the cost of electricity been rising but consumers are finding limits to the infrastructure that delivers power to their centers. Capturing cost savings produced by consolidation of equipment to a central site must be scrutinized in relation to power available within that site. Since the DMX-4 series intentionally offers consolidation savings, it has been engineered with power efficiency in mind. The following improvements in the power arena are available with the DMX-4:

- **Performance and scalability:** Improved caching algorithms allow consolidation of workloads. Symmetrix DMX-4 customers can get the same performance with less hardware, reducing energy requirements.
- **New tiered storage options:** Both the Symmetrix DMX-3 and DMX-4 platforms offer the industry's broadest range of energy-efficient disk drives, which provide higher capacity for more terabytes per square foot. The Symmetrix DMX-4 is even more efficient through the support for 500 GB and eventual 750 GB or larger SATA II drives (750 GB is targeted for the end of 2007). Figure 6 shows a drive power comparison.

750 GB SATA II Data on Different Capacity/Performance Drives Lower Capacity Drives Consume More Power per Terabyte




Figure 6. Drive power comparison

- **EMC Power Calculator:** Version 2.0 allows the user to make accurate estimates of energy use for specific configurations. This can be especially valuable in planning changes and consolidations. The calculator includes previous generations of Symmetrix, CLARiiON®, and Celerra®, as well as the current version of all EMC platforms. It reports on operational cost, weight, energy consumption, sound, and floor space. The calculator has number of ease-of-use features such as the ability to save configurations and reports.




Figure 7. EMC Power Calculator

The DMX-4 series benefits from a planning tool and tiering options that allow users to specifically tailor DMX-4 storage arrays to meet current and predicted energy limitations. The Enginuity Storage Operating Environment contains improvements to make the DMX-4 as energy efficient as possible.

Conclusion

EMC continues to deliver balanced solutions for information and storage management. Technology refreshes work in partnership with software updates to create products that address traditional business needs as well as the cutting-edge requirements of new business paradigms. Enginuity 5772 and the Symmetrix DMX-4 series are a fine example of this balance. Enginuity 5772 and the DMX-4 series are a marriage of hardware evolution with feature enhancements. Working together, new hardware and the Enginuity Storage Operating Environment deliver the next step in high-end storage array productivity. Business continuance, disaster recovery, drive capacity, efficiency of cache, RAID XOR optimization, and 4 Gb/s bandwidth with point-to-point drive carriers all blend into a powerful, productive tool ready to deliver advanced storage functionality for today's business processing needs. Browser-based management software simplifying common tasks and a purposeful focus on power efficiency demonstrate an EMC commitment to meeting the latest customer challenges as well as the traditional business concerns. The DMX-4 system justifiably extends EMC's leadership in the high-end enterprise storage market.