


EMC Symmetrix DMX Architecture

Product Description Guide

Defining the Next Generation in High-End Storage Systems





Advanced architecture • Landmark performance • Unmatched availability • Richest functionality • Unexpected economics

Table of Contents

Chapter One: Introduction	5
Factors driving the need for high-end storage	5
High-end storage requirements	6
Sustained, burst, and business-continuity performance	6
Availability	7
Functionality	7
Economics	7
More exacting high-end storage requirements test the limits of existing storage architectures	8
Chapter Two: The Symmetrix DMX Architecture	9
Emerging high-end storage requirements	9
The direct matrix architecture versus switch and bus-based architectures	9
Symmetrix: The standard for high-end storage	9
Symmetrix foundation—hardware and software	10
Shared-model design	10
Engenuity storage operating environment	10
DMX – The next-generation Symmetrix architecture	11
DMX Messaging Matrix	14
DMX Communications Control and Environmental Matrix	16
Communication Control Module	16
Environmental Control Module	17
Symmetrix DMX: a “future-ready” storage architecture	18
The Symmetrix DMX Architecture meets both criteria	19
Summary	19
Chapter Three: Performance	21
Symmetrix DMX: Uncompromising Performance	21
Symmetrix DMX Scaling Attributes	22
Configuring to the environment	23
Symmetrix DMX Consolidation Benefits	23
Symmetrix DMX Burst Performance	24
The role of Engenuity in burst performance	24
The role of the Symmetrix DMX Architecture in burst performance	24
Business Continuation Performance	25
Summary	26
Chapter Four: Availability	27
Symmetrix: The gold standard in high-end storage	27
DMX Control and Environmental Matrix	27
Communication Control Module	28
Environmental Control Module	28
Non-Disruptive Operations	29
Symmetrix DMX Error Detection and Remote Support	30
Symmetrix component-level redundancy	31
Global Memory Director redundancy	31
Memory chip redundancy	32
Triple Modular Redundancy with Majority Voting	32
Channel Director redundancy	33
EMC PowerPath	33

TABLE OF CONTENTS

Internal control data path redundancy	33
Power supply redundancy	34
Battery backup	35
Fibre Channel back-end redundancy	35
Arbitrated loop design	35
Symmetrix Component-Level Serviceability	36
Online Enginuity upgrades	37
Symmetrix Data Protection Options	38
Symmetrix DMX Mirrored and Parity RAID Options	38
Disk mirroring concepts	38
Parity RAID (3+1) and Parity RAID (7+1)	39
RAID 5 (3+1) and RAID 5 (7+1)	39
Summary	40
Chapter Five: Symmetrix DMX Functionality	41
Enginuity	41
Foundation	41
Performance	41
Availability	41
Openness	42
Enginuity Foundation	42
Security: Access Controls	43
Enhanced System Audit and Investigation	43
Enginuity Performance: Dynamically Adjusting Performance Algorithms	43
EMC Enginuity Performance Solutions	44
Maximizing Global Memory Performance	44
Intelligent Adaptive Pre-fetch Algorithm	44
Quality-of-Service (QoS) Algorithms	45
Permacache	45
Maximizing Disk Performance	45
Dynamic Mirror Service Policy (DMSP)	46
Rotational Position Ordering (RPO)	46
Enginuity Availability	46
Continuous integrity testing and error detection	47
Parity Check	47
Checksums and Cyclical Redundancy Checks	47
Continuous trending analysis	48
Failure containment and recovery	48
Proactive Cache and Disk Scrubbing	48
Checking the checkers	48
Instruction monitoring	48
Non-Disruptive Upgrades (NDUs)	49
Change Management	49
Enginuity's Open Environment	49
Open Standards	49
Networked Storage	50
Interoperability	51
Networked Storage Evolution: Driven by Consolidation	51
Storage Area Networks	52
Cost reduction through consolidation	52
Backup and recovery	52

SAN in a Symmetrix DMX environment	53
Networked Attached Storage	53
Integrated SAN and NAS	53
EMC Business Continuance Software	54
Alternative solutions for long-distance replication	56
The Symmetrix Direct Matrix Architecture software performance	58
Open Storage Management	58
Intelligent Supervision	59
Device Management	60
Information Safety	60
Infrastructure Services	61
Summary	62
Chapter Six: The Economics of Symmetrix DMX	63
A paradigm shift in the economics of high end storage	63
Modular flexibility: a key attribute of the Symmetrix DMX Architecture	63
Symmetrix DMX series systems	64
Symmetrix DMX800	64
Symmetrix DMX1000	65
Symmetrix DMX2000	65
Symmetrix DMX3000	65
Symmetrix DMX optional configurations	66
The economics of acquisition	66
The economics of consolidation	67
The economics of RAID 5 and Parity RAID	67
Summary	68
Chapter Seven: Services	69
EMC Services	69
Information Solutions Consulting: Evaluating Change	69
EMC Services: Planning, Implementing, and Initiating Operations	70
Proven Methodologies	70
Customer Service: Ongoing Operations	71
Remote Support	71
Software Support	71
Change Control	71
Installation Support	71
Post-sale Warranty and Product Support	71
Worldwide Organization, Local Support	71
Global Technical Training	72
The EMC Proven Professional Certification Program	72
EMC Powerlink	72
Chapter Eight: Summary	73
Summary	73
Symmetrix DMX and Automated Networked Storage	73

Chapter One: Introduction

Factors driving the need for high-end storage

In most industries there are business-critical applications or sets of applications whose service levels transcend the others by a wide margin. The pressure on IT organizations tasked with providing these service levels is intense, and will become more so over time.

Many factors contribute to this inexorably mounting pressure:

- Contractually-binding service level agreements (SLAs) require IT organizations to deliver agreed-upon performance, availability, and other measurable support metrics. These metrics must be met under all circumstances, without fail. For example, guaranteed up-time and access speed might be stipulated for certain information and applications during typical load periods as well as during unexpected spikes in activity. These SLAs remain constant over time as application data requirements and demands continuously increase, resulting in additional strain on staff and technical infrastructure.
- The world is experiencing an information explosion. Since the beginning of time, the human race has created 17 billion gigabytes of information. According to a UC Berkeley study, the amount of new information being produced each year is between one and two exabytes. (An exabyte is a number followed by 18 zeros. Hence, one exabyte = 1,000,000,000,000,000,000 bytes). A higher volume of total information inevitably drives transactional workloads for IT higher as more and more users demand access to a huge and steadily growing mass of data.
- Recent manmade and natural disasters have created additional demand for information storage solutions that protect data and ensure that businesses with high-end storage requirements can withstand even the most dire catastrophe with minimal data loss.
- In accordance with Moore's Law (promulgated in its original form by Intel co-founder Dr. Gordon E. Moore in 1964), the volume of information that can be accommodated by a given amount of silicon is roughly doubling every 18 months. At the same time, CPU speeds are doubling every year, and interconnect bandwidth is projected to grow by a factor of ten over the next five years. The result: far greater workloads for IT.
- Regulatory requirements such as those contained in Securities and Exchange Commission Rule 17a-4, Federal Food and Drug Administration 21CFR Part 11, Health Insurance Portability and Accountability Act (HIPAA), Sarbanes-Oxley Act, and Department of Defense 5015.2, are also driving storage demands on the high end.

The non-negotiable requirement to satisfy constantly rising service levels in the face of uncompromising SLAs, steadily increasing transactional demand, fears of catastrophic data loss, order of magnitude leaps in CPU speeds and bandwidth, and stringent new regulatory requirements severely stresses the IT infrastructure. And in no component of the environment is the pressure felt more keenly than in storage.

Storage plays a vital role in delivering IT services, directly impacting the performance, availability and functionality of most applications and, by extension, the ability of the enterprise to fulfill its mission.

High-end storage requirements

All computing environments require some degree of performance, availability, and functionality. What distinguishes the high-end from the mid-tier is that the former requires the highest possible levels of performance, availability and functionality, without compromises or trade-offs. For example, settling for less performance in order to add more functionality in the form of new applications is not an option.

Sustained, burst, and business-continuity performance

The high-end storage environment requires the best possible performance both in routine operations (sustained performance), and during sudden, dramatic spikes in I/O demand (burst performance).

The phrase “routine operations” needn’t imply a sedate daily regimen. A large market information research company might routinely process hundreds of millions of checkout scanner records per week, and maintain hundreds of billions of records at any given time. It might be routine to aggregate this information into extremely large data structures, and then break it out for clients who insist upon customized views of the data, requiring extremely high levels of performance and sustained bandwidth to complete the data analysis.

In a high-end environment, sustained operations must be unaffected by a sudden spike in I/O, such as an unanticipated full database query during a batch run. The ability to deftly accommodate such demands, without trepidation or hesitation, could be key to quickly seizing a business opportunity, or to avoiding a costly mistake.

In addition, high performance in business continuance tasks such as remotely replicating and restoring data is emerging as a bottom-line requirement. If, for example, a large retail company should lose a database, the underlying storage infrastructure must recover quickly. Failure to do so might result in millions of dollars in losses, and possibly even imperil the survival of the business.

As in all aspects of high-end storage, performance trade-offs are not an option. Sustained, burst, and business-continuity performance must all be maintained at the highest levels at all times – no matter what.

Availability

At the high-end, brief periods of system downtime or data unavailability caused by either unplanned or planned outages that might be tolerable (or at least survivable) in less demanding environments can have devastating consequences. Beyond loss of revenue, the impact in missed business opportunity and damaged customer relationships can be incalculable. The possibility of any such eventuality must be minimized to the extent humanly possible.

Functionality

The IT infrastructure is nothing more than a delivery mechanism for software-based functionality, wherein resides the actual business value to the enterprise.

The word “functionality” encompasses all of the applications, software products, and operating environment capabilities that 1] manipulate, move and store information, and 2] implement and enhance performance and availability of IT resources. For the organization facing escalating competitive pressures and rising service level demands in a high stakes business environment, functionality is of paramount importance.

Economics

The challenge of providing uncompromising high-end performance, availability, and functionality has been complicated in recent years by a precipitous decline in IT budgets and resources. Most IT organizations are now required by their management to do far more than ever before with much less than ever before. As a result, vastly improved economics of acquisition and ownership have emerged as yet another high-end storage requirement.

More exacting high-end storage requirements test the limits of existing storage architectures

Existing storage architectures do a competent job at providing performance, availability, functionality, and reasonable economics. The challenge is that all of the factors driving up service levels are not only unabated, but are intensifying. SLAs are more demanding than ever; information volume continues to increase transaction demand; high availability requirements continue to challenge organizations both financially and operationally, and regulatory compliance places more demands on information lifecycle management. In short, it is becoming increasingly difficult to handle present high-end storage challenges, let alone the even tougher ones looming closer every day.

The quest to meet both present and emerging high-end storage challenges led EMC engineers to develop a storage architecture that is a quantum leap beyond existing technologies. It delivers unprecedented performance, availability, and functionality, while at the same time effecting a paradigm shift in the economics of high-end storage. This architecture, Symmetrix DMX™, is introduced in Chapter Two.

Chapter Two: The Symmetrix DMX Architecture

Emerging high-end storage requirements

As discussed in Chapter One, IT organizations today must deliver higher service levels for performance, availability, and functionality, with fewer people and smaller budgets, at a time when numerous factors are driving IT workloads to unprecedented levels.

In response to this challenge, EMC engineers designed a direct matrix high-end storage architecture, Symmetrix DMX.

Symmetrix DMX delivers more performance, higher availability, and more functionality, with vastly improved economics. It is the foundation for the next decade of high-end Symmetrix® storage arrays. The design is capable of keeping up with the most aggressive service levels even as CPU speeds, network bandwidth, and disk capacities continue to grow exponentially.

The direct matrix architecture versus switch and bus-based architectures

Switch and bus-based designs offer good performance at an attractive price, but their ability to scale in very demanding environments is limited.

Even more significantly, busses and switches present issues with contention, latency, bandwidth and potential points of failure – all of which can adversely impact service levels.





The direct matrix design of Symmetrix DMX represents a significant improvement over switch and bus-based architectures.

The underlying architectural evolution driving the direct matrix concept is in the word “direct.” Direct connections within a storage system are better for precisely the same reason that a direct airline flight from point A to point B is better. A direct airline connection eliminates any possibility of delay at an intermediate airport. Similarly, the Symmetrix Direct Matrix Architecture™ delivers direct access from the front of the storage array to the back, guaranteeing the highest possible I/O throughput. There is no possibility of delay at a bus or switch or hub. Data flow is unimpeded.

Symmetrix: The standard for high-end storage

The Symmetrix DMX Architecture carries forward many of the architectural attributes that have made Symmetrix the standard storage architecture for mission-critical computing environments.

Symmetrix DMX Architecture core components are: Channel Directors for host communication; Disk Directors for disk communication; and Global Memory Directors for I/O delivery from hosts to Disk Directors.



Symmetrix core components are: Channel Directors for host communication; Disk Directors for disk communication; and Cache Directors for I/O delivery from host to disk Directors.

Symmetrix foundation – hardware and software

Shared-model design

The synergy between the Symmetrix hardware and software architectures has made Symmetrix systems best-in-class for the last decade.

The design point for any high-end storage array is to move I/O from server to disk and disk to server in the most efficient manner possible. All Symmetrix storage arrays, past and present, accomplish this by means of a shared-model architecture. In fact, the first Symmetrix introduced in the market was a symmetric multi-processing (SMP) shared-model design.

The shared-model design means that all directors (channel and disk) are autonomous and work in parallel. It also means that each director, with up to eight PowerPC processors, can access all regions in the Global Memory Directors.

Because access to the global memory resources in the system is symmetrical, any processor on a director can participate in any event within the system. Examples of such events include (but are not limited to) I/O READS and WRITES, error detection and correction, and creation of remote and local copies of data. The directors can also provide resilience if a fault occurs in another specified director. This “shared everything” design is key to delivering consistently high service levels for performance, availability, and functionality.

Enginuity storage operating environment

Although the “shared everything” hardware design allows all components to share the workload, it is the Enginuity storage operating environment that provides intelligence to Symmetrix systems.

Enginuity is an event-driven storage operating environment that prioritizes multiple simultaneous events within the system and guarantees quality of service (QoS) for the most important events. For example, Enginuity ensures that correcting a two-bit error in memory takes priority over a WRITE request, a READ request takes priority over updating a business continuance volume (BCV), and so on.

Enginuity also delivers performance, data integrity, and open integration for the Symmetrix family of storage arrays. (See Chapter 5 for more detail on the Enginuity storage operating environment.)

Symmetrix: A massively parallel storage system designed to multi-task numerous simultaneous events

The combination of the Symmetrix multi-processing hardware architecture with the event-driven Enginuity storage operating environment produces a massively parallel storage system designed to multi-task numerous simultaneous events.

For example, when a new WRITE I/O is committed to memory, the new data is immediately available to all of the CPUs on every director board. While the data is protected in memory, the processors on all of the director boards can work autonomously on the new data to update a mirrored pair; send the update over an SRDF® link; update a business continuance volume; report the current status of all events to the management software; and handle error detection and correction of a failed component.

All of these tasks can occur simultaneously, without de-staging to disk and re-staging to a separate region in memory. This sophisticated functionality makes Symmetrix storage arrays best-in-class for high-end storage requirements.

DMX – The next-generation Symmetrix architecture

The core of the previous generation Symmetrix systems is carried over to the Symmetrix DMX storage arrays. However, the two major design points for the Symmetrix DMX architecture—the Direct Matrix Interconnect and the Global Memory Architecture—more fully leverage the inherent capabilities of the massively parallel hardware design and event-driven storage operating environment. As a result, Symmetrix DMX systems deliver significantly greater performance, availability, and functionality than previous Symmetrix systems and all existing competitive systems.


The Direct Matrix Interconnect is a matrix of dedicated high-speed links to all of the system components, providing unparalleled internal aggregate bandwidth of up to 64GB/sec.


This underlying interconnect matrix allows the architecture to deliver four times the bandwidth, five times the processing power, and 10 times the memory throughput of the nearest competitor.

The Global Memory Director manages 32 independent global memory regions. It essentially eliminates contention for shared global memory while optimizing utilization of system resources.

A layout of a Global Memory Director is shown below.


The Symmetrix DMX global memory architecture provides unprecedented dedicated paths for concurrent cache operations.





The Symmetrix DMX Direct Matrix Interconnect architecture provides direct access to all (up to 32) memory regions simultaneously. (See diagram below.)

The architecture also allows great flexibility in scaling up and down to meet the specific capacity and connectivity needs of a given environment, bringing unexpected economics to high-end storage.



The Symmetrix DMX matrix with 128 direct point-to-point connections and 32 memory regions provides a data matrix rated to 64 GB/s of internal aggregate bandwidth supporting up to 32 concurrent global memory operations.

Each of the point-to-point connections between director and global memory are composed of several serial links. Each pathway from director to memory is a two-bit-wide connection running at 1.25 billion transfers per second. Likewise, the pathway from memory to director is another two-bit-wide connection, also running at 1.25 billion transfers per second. This makes the raw bandwidth between director and global memory $4 \text{ bits} \times 1.25 \text{ GB/s} = 5 \text{ GB/s}$.

This raw bandwidth is encoded with an industry-standard 8b/10b, so for every ten bits sent, eight can be counted as data. Thus, the data bandwidth of the connection is $5 \text{ GB/s} \times 8\text{b}/10\text{b} = 4 \text{ GB/s}$ or 500 MB/s.

The 8b/10b encoding provides data protection on the channel against errors. It also provides in-band control of the channel, and allows for end-to-end synchronization of operational status. For additional data integrity the data is sent in 512 byte packets with each packet protected by two additional 16-bit Cyclical Redundancy Checks (CRCs)—one for the header and one for the data.

Additionally, the entire data set of the I/O is protected by a separate high-level CRC in the Global Memory Directors. This provides additional protection over the channel at three independent levels: link, transport, and logical.

The combination of the Direct Matrix Interconnect and the Global Memory Architecture has redefined high-end storage capabilities for performance, availability and functionality.

DMX Messaging Matrix


The Symmetrix DMX inter-processor control and communications mechanism has been completely redesigned to take advantage of the latest state-of-the-art interconnect technology. The Symmetrix DMX Message Matrix consists of point-to-point interconnects between all directors within the system. It has been designed to provide high bandwidth and low latency control communication, which manages data movement throughout the Symmetrix system. The Message Matrix expedites all event requests (e.g., READ/WRITE requests, Symmetrix Remote Data Facility commands, business continuance volume updates) throughout the system to deliver unprecedented performance.

In previous Symmetrix systems, as well as in competitive offerings, control and messaging communication is handled through “mailboxes” in memory. With mailbox technology, requests and notifications for operations are stored in the target mailbox by the requesting Director, and target Directors asynchronously scan the mailboxes in memory to see if they have tasks to process (e.g., READ/Writes, remote replication, local replication). This implementation is effective, but comes with a cost in efficiency due to the CPU overhead and the contention for global memory access.


In the Symmetrix DMX Messaging Matrix, this processing overhead is eliminated, as is contention for the global memory resources. This frees the processors to do more I/O processing and enables more of the global memory bandwidth to be utilized for moving data instead of controlling information.

Serial I/O connections versus parallel buses for high speed data transmission

Parallel bus architectures have been a mainstay in computer designs for many years. Composed of many wires (and, therefore, multiple connectors, drivers, and receiver chips), a bus is typically designed to be shared among several elements. Performance of a shared-bus implementation is limited. As performance demand increases, each unit attached to the bus requires more transactions. Wait times caused by contention on the bus begin to increase, even as bus cycle times are reduced. A point-to-point, serial I/O connection scheme, as implemented in the Symmetrix DMX Architecture, eliminates the resource contention associated with the shared bus design. In addition to superior performance, it delivers superior interconnect scalability and data integrity.



The Symmetrix DMX fully redundant messaging matrix provides high-bandwidth and low-latency control communication to manage data movement throughout the Symmetrix system.



Redundancy in the Message Matrix is a fundamental characteristic of the high-availability of the Symmetrix DMX hardware.

Independent, redundant fault domains are intelligent at the hardware level and balance load traffic between each domain. If a fault occurs within one of the domains, the traffic is seamlessly routed by the hardware to the available domain and the system is alerted to the faulty domain, which can then be replaced.

DMX Communications Control and Environmental Matrix

Remote Support is an important and integral part of EMC Customer Service and Support.

Every Symmetrix unit has an integrated service processor that continuously monitors the Symmetrix system. The service processor can communicate with the EMC Customer Support Center through a customer-supplied direct phone line. The service processor automatically dials the Customer Support Center when the Symmetrix system detects a current or potential component failure or environmental violation. An EMC Product Support Engineer at the Customer Support Center can also run diagnostics remotely via the service processor to determine the source of a problem and potentially resolve it before the problem becomes critical.

Within the DMX control matrix are the Environmental Control Module (ECM) and the Communication Control Module (CCM).

These two modules provide the low-level system-wide communications for running application software, monitoring, and system diagnostics from the service processor.


In the event of a module failure, the Communication Control Module and Environmental Control Module can be replaced non-disruptively by utilizing redundant module functions within the Symmetrix DMX Architecture.

Communication Control Module

The Communication Control Module's primary function is to act as a communications agent between the laptop service processor and the embedded processing nodes within the system.

External connections to the service processor provide dial-home capabilities for remote monitoring and diagnostics.

A software driven failover mechanism ensures that the Ethernet remains intact should a CCM become temporarily unavailable.



Redundant communication control modules act as a communications agent between the embedded processing nodes, as well as executing commands from the service processor.

As a backup to the Ethernet fabric, the CCM also contains an alternate means of serial communication to embedded processing nodes within the DMX system. The primary service processor can communicate to individual nodes should the Ethernet become unavailable.

The CCM can also issue remote commands to Channel, Disk, and Global Memory Directors and the Environmental Control Module. These commands can be issued from the laptop service processor or remotely by the EMC Customer Support Center.

Also within the CCM is the top-level fabric for the Message Matrix communications system. This is a high-speed communications fabric within the Symmetrix DMX Architecture that allows for fast, reliable messaging between compute nodes.

Environmental Control Module


The Environmental Control Module (ECM) monitors and logs environmental events for all Symmetrix DMX field replaceable units (FRUs) and reports any operational problems, such as thermal excursion, voltage drops, etc. FRUs include Channel and Disk Directors, Global Memory Directors, port bypass cards, power supplies, battery backup units, thermal monitors, fans, disk drives and switches.

Symmetrix DMX: a “future-ready” storage architecture

A point frequently reinforced in this document is that service levels and the demands they impose upon IT organizations are steadily escalating with no end in sight. The future will, without a doubt, be far more challenging than the present.

Anticipating this reality, it is essential that a high-end storage architecture be future-ready, meaning it can:

- without modification, handle far greater demands than are presently placed upon it
- incorporate future technological enhancements to address future service level requirements (including some that might be far beyond those that can be currently foreseen).



The Environmental Control Module (ECM) monitors and logs environmental events for all Symmetrix DMX field-replaceable units (FRUs) and reports any operational problems, such as thermal excursion, voltage drops, etc.

The Symmetrix DMX Architecture meets both criteria

The Direct Matrix Interconnect and DMX Global Memory Architecture were specifically designed to accommodate exponentially escalating interconnect bandwidth and global memory throughput demands.

And, by design, the architecture can absorb and leverage new processor technologies, interconnect protocols, and storage media designs as they evolve. Using today's available technology, it is capable of supporting more than 2,048 disk drives (a four-fold increase). In the future, the DMX architecture will allow EMC to easily implement such performance enhancements as 512GB of global memory, 128 global memory regions, 10Gb Fibre Channel disk drives, or 2GHz processors.

As a result, the Symmetrix DMX Architecture (and IT organizations that invest in Symmetrix DMX systems) can take the future in stride.

Summary

Emerging high-end storage requirements are driving the need for unprecedented levels of performance, availability and functionality at a time when IT budgets are being slashed. The Symmetrix DMX Architecture is EMC's response to this challenge.

The Direct Matrix design eliminates contention and scaling issues associated with switch and bus-based architectures. The architecture delivers direct access from the front of the storage array to the back, ensuring the highest possible I/O throughput.

The Symmetrix Direct Matrix Architecture carries forth the essential characteristics that have made Symmetrix systems the standard for high-end storage in the past decade. The Symmetrix shared-model design moves I/O from server to disk in the most efficient possible manner.

The Enginuity storage operating environment prioritizes multiple simultaneous events within the system and guarantees quality of service (QoS) for the most important events. Enginuity provides the intelligence to derive maximum functionality from the Symmetrix architecture.

The two major design points for the Symmetrix DMX Architecture are the Direct Matrix Interconnect and the Global Memory Architecture. These innovations more fully leverage the inherent capabilities of the massively parallel hardware design and event-driven storage operating environment.

The Direct Matrix Interconnect is a matrix of dedicated high-speed links to all of the system components, providing internal aggregate bandwidth of up to 64GB/s. It allows the architecture to deliver four times the bandwidth, five times the processing power, and 10 times the memory throughput of the nearest competitor. The Global Memory Director manages 32 independent global memory regions, eliminating contention for shared global memory while optimizing utilization of system resources. The combination of the Direct Matrix Interconnect and the Global Memory architecture has redefined high-end storage capabilities for performance, availability, and functionality.

The Symmetrix DMX Architecture is “future ready.” It can, without modification, handle far greater demands than are presently placed upon it, and readily incorporate future technology enhancements as service level requirements continue to rise.

Chapter Three: Performance

Symmetrix DMX: Uncompromising Performance

The Symmetrix DMX Architecture has an enormous capacity for I/O throughput. This translates into performance capabilities that far exceed any bus or switch-based products on the market.

The superiority of the Symmetrix DMX Architecture is strikingly apparent in applications that require very high levels of sustained performance and for environments that experience frequent “bursts” of I/O activity and demand.

Among the benefits of elevating the number of I/O operations per second (IOPS) that a system can deliver is that more online transactions (OLTP) can be performed per unit of time, improving response times and increasing the efficiency of revenue generation.

Another benefit of elevating the number of IOPS is that decision support systems (DSS) can be loaded faster which correlates to fresher or more timely data to support critical business decisions.

However, in a dynamic IT environment with high-end storage requirements, where sudden bursts of I/O are common, sustained performance is not enough. The ability to handle sudden, heavy workload surges, without degrading application response times or imperiling service levels, is a bottom-line requirement. One of the most valuable characteristics of the DMX Architecture is its ability to maintain the highest sustained OLTP and DSS performance even during sudden bursts of I/O activity.

Of equal importance to its sustained and burst performance capabilities, the DMX Architecture provides an exponential boost in business continuance performance. Local and remote copies of data can be made faster and more often than ever before. This protects the enterprise from the potentially devastating impact of losing substantial amounts of business-critical information, and facilitates a return to normal operations when problems occur.

Delivering uncompromising sustained, burst, and business continuance performance today is a major challenge. And, the challenge can only intensify as CPU speeds double every 18 months, interconnect bandwidth increases by a factor of ten over the next five years, and the number of users requiring access to data continues to grow. The Symmetrix DMX Architecture is the only available architecture that is equal to this challenge.

Sustained I/O performance of Symmetrix DMX systems is not compromised even if a controller is lost, workloads spike, or while synchronizing, splitting, or restoring full copies of production data.


In short, no matter what happens, applications are unlikely to suffer any adverse impact whatsoever, dramatically increasing the probability that SLAs for performance will be met at all times.

Symmetrix DMX Scaling Attributes

The unlimited scalability of the Symmetrix DMX Architecture is among its most revolutionary attributes.

The architecture's flexibility enables an IT organization to select the Symmetrix DMX system that closely matches its high-end storage needs in the most economical package. (See Chapter Six.) The flexibility of the architecture also facilitates customized configuration of the system to match very specific capacity, connectivity, and bandwidth requirements.


A unique characteristic of the Symmetrix DMX Architecture is that connectivity and bandwidth are increased as components are added to the system. In contrast, adding components to bus and switch-based architectures exacerbates contention issues, adversely impacting performance.



Bus and switch-based architectures also lack the ability to scale up and down economically. This forces IT organizations to resort to sub-optimal and costly expedients. For example, monolithic systems might be under-configured to accommodate smaller capacity and connectivity needs while still trying to maintain high-end performance and availability. The enterprise ends up over-paying in infrastructure costs (including power and cooling) for a solution that, in the end, might not closely correlate with performance and availability requirements.

The statement above, that connectivity and bandwidth in the Symmetrix DMX Architecture increase as components are added to the system, should not be interpreted to mean that a rack-mounted Symmetrix DMX800 configuration with fewer components (two Channel, two Disk and two Global Memory Directors) delivers exponentially less performance than a larger Symmetrix DMX system. In this case, the performance gap is substantially narrowed by an innovative approach to connectivity, as shown below.

In the Symmetrix DMX800, all eight of the 1.2GB/s channels on each Channel and Disk Director are connected through global memory. As a result, the smallest Symmetrix DMX system achieves better performance than the closest competitor with fewer components and more attractive economics.



The modular design of the DMX Architecture can be tailored to provide front- and back-end configurations to meet specific host requirements and application performance.


Configuring to the environment

With the Symmetrix DMX Architecture, front-end or back-end configurations can be defined with great flexibility. As a result, the IT organization can acquire only the exact amount of storage needed, configured to the precise requirements of the environment.

The following are examples of tailoring a Symmetrix DMX configuration to the specific needs of the environment.

Applications that are latency sensitive, such as OLTP, benefit from many host connections—the more the better—so a configuration for this environment would feature more front-end directors. (See diagram below.)

The Symmetrix DMX Architecture scales front-end aggregate bandwidth linearly with models that have more Channel Directors.




Very demanding bandwidth sensitive applications, such as data strip-mining, prefer greater concurrent disk activity, hence more back end directors as shown in the following diagram.

As the above examples illustrate, the flexibility of the Symmetrix DMX Architecture enables IT organizations to acquire the exact amount of storage needed, configure it to the precise requirements of the environment, and then scale up incrementally on a “buy only what’s necessary, when necessary” basis.

Symmetrix DMX Consolidation Benefits

The Symmetrix DMX architecture facilitates consolidation of servers and storage. Literally thousands of servers can be consolidated into a single array with dozens of terabytes of capacity, while maintaining high-end service levels for performance, availability and functionality. The result: the most efficient possible use of resources and lowered operational costs.



Targeted configurations with additional front-end directors can be optimized to support high server counts in large storage consolidation initiatives.

Symmetrix DMX Burst Performance

Sudden, heavy workload spikes and surges are an unavoidable reality in all IT environments. The challenge is to accommodate extraordinary, unexpected I/O requirements, such as an unanticipated full database query during a batch run, without compromising normal service levels.

The ability of Symmetrix DMX to deliver unprecedented performance even during unexpected bursts of I/O activity is a function of 1] the Enginuity storage operating environment and 2] the modularity of the architecture itself.

The role of Enginuity in burst performance

Competitive architectures mindlessly respond to requests in the order they are received – the worst possible approach to a workload surge, and one that is guaranteed to adversely impact service levels.

In contrast, Enginuity intelligently mobilizes system resources to deal with the surge in the most efficient and nondisruptive possible manner. (This is a key advantage of a mature, sophisticated storage operating environment. For more on Enginuity, see Chapter Five.)

Enginuity's ability to prioritize and allocate is, in part, innate—the product of twelve years of accumulated algorithmic intelligence. It is also guided by user-defined quality-of-service settings for logical volumes and applications. Enginuity knows how to handle large amounts of extra work, without compromising normal service levels.

The role of the Symmetrix DMX Architecture in burst performance

Intelligent allocation of resources is futile, unless the resources themselves are capable of a flexible, high-performance response. Here, once again, the modularity and enormous number of direct high-speed connections of the Symmetrix DMX Architecture dramatically enhance performance. There is no bus or switch contention to create bottlenecks.

Processor contention is also eliminated, because all processors can share the load.

For example, if a burst of random I/O READS comes into the system from a single server or application (assuming it is attached to two front-end Channel Directors) it will 1] be handled by eight PowerPC processors 2] have 16 500MB/s channels through Global Memory, and 3] once again be handled by eight CPUs on the Disk Directors.

Enginuity, the event driven Symmetrix storage operating environment, can prioritize tasks across all of the available channels and processors.

For example, if a surge of I/O READ requests hits the system while one processor on a Disk Director is updating a business continuance volume (BCV), Enginuity will find the processor that will most efficiently service the READ request and move the task of updating the BCV to one of the other processors on the associated Disk Director through the shared global memory.

Although all previous Symmetrix systems were capable of preemptive multi-tasking, the additional surge of I/O caused some degree of contention because of the limited bus resource.

The closest competitive architecture also has the issue of contention, and is even further handicapped in comparison with Symmetrix by cache, operating system, and hardware limitations, as outlined below.

COMPETITIVE SYSTEM VS. **SYMMETRIX**

- four discrete caches with four separate mailbox lookup tables vs. **global memory**
- control-based microcode vs. **an event-driven, intelligent storage operating environment**
- a task-isolated controller vs. **a shared multi-tasking hardware platform**

The bottom line: the ability of Symmetrix DMX systems to handle burst performance far surpasses any previous Symmetrix system or competitive offering.

Business Continuity Performance

The unexpected burst of I/O is not the only type of workload surge common to all high-end environments. Remote replication, full volume copying, full volume restores, and other business continuity measures are also daily realities of high-end storage management. (For details on Symmetrix business continuity software, see Chapter Five.)

These activities take time, and are usually done in off-hours to avoid compromising normal production. As a result, replication might occur as infrequently as once every twenty-four hours. In a demanding, dynamic environment, the problem is obvious. In a worst-case data-loss scenario, the enterprise might find itself relying on replicated data that is up to 23 hours old. For a retail chain on December 23rd, this would be disastrous.

Again, the reason replication and copying is not done more often is that it takes too long and impacts routine performance—constraints that do not apply to Symmetrix DMX.

Among the most revolutionary capabilities of the Symmetrix DMX Architecture is the speed with which data can be copied, restored, and moved to various locations. (As discussed above relative to burst performance, these performance capabilities are a function of the Enginuity storage operating environment and the Direct Matrix Architecture.)

As an example, creating a business continuity volume to load a data warehouse can be accomplished three to ten times faster with Symmetrix DMX. The IT administrator who has a shrinking window of time in which to do database backups can now replicate three to ten times more data in the same amount of time. In short, more information can be backed up more often, and over longer distances.

Summary

The Symmetrix Direct Matrix Architecture delivers sustained, burst, and business continuance performance superior to that of previous Symmetrix systems and all existing competitive systems. The highest levels of sustained performance are never compromised by unexpected workload spikes or business continuance tasks.

The uniquely flexible Symmetrix DMX Architecture easily scales up and down, and systems can be custom-figured to meet the precise needs of the environment. IT organizations can acquire only the storage they need, configure it to their requirements, and scale up on a “buy only what’s necessary, when necessary” basis.

The Symmetrix DMX Architecture, with its sophisticated Enginuity storage operating environment, is equal to the exacting service-level challenges of today, and the even tougher challenges of tomorrow.

Chapter Four: Availability

Symmetrix: The gold standard in high-end storage

As service levels for critical applications escalate exponentially over the next decade, so will requirements for information availability and data integrity.


Symmetrix is the gold standard for mission-critical applications. It has proven itself time and again, over twelve years, in the world's most demanding environments, including the data centers of the largest financial, insurance, and telecommunications companies.

Symmetrix was engineered to work flawlessly, to continue to run no matter what, and to be serviced proactively and non-disruptively.

Symmetrix DMX now raises the availability bar even higher with the world's most advanced fault-tolerant design featuring full redundancy, proactive monitoring, and error detection and correction.

DMX Control and Environmental Matrix

Remote support is an important and integral part of EMC Customer Service and Support. Every Symmetrix unit has an integrated service processor that continuously monitors the Symmetrix environment. The service processor can communicate with the EMC Customer Support Center through a customer-supplied direct phone line. The service processor automatically dials the Customer Support Center when the Symmetrix system detects a component failure or environmental violation. An EMC Product Support Engineer at the Customer Support Center can also run diagnostics remotely via the service processor to determine the source of a problem and potentially resolve it before the problem becomes critical.



Within the DMX control e-net matrix are the Environmental Control Module (ECM) and Communication Control Module (CCM). These two modules provide the low-level systemwide communications for running application software, monitoring, and system diagnostics from the service processor.

The Communication Control Module (CCM) and Environmental Control Module (ECM) feature non-disruptive replacement of failed modules by utilizing redundant module functions within the DMX Architecture.

Communication Control Module

The Communication Control Module's primary function is to act as a communications agent between the laptop service processor and the embedded processing nodes within the system. External connections to the service processor provide dial-home capabilities for remote monitoring and diagnostics.

A software-driven fail-over mechanism ensures that the Ethernet fabric remains intact should a CCM become temporarily unavailable.

As a backup to the Ethernet fabric, the CCM also contains RS232 multiplexing logic to allow for an alternate means of serial communication to embedded processing nodes within the DMX system. The primary service processor can use its RS232 serial port to communicate to individual nodes should the Ethernet fabric become unavailable. The CCM can also issue remote commands to the director boards, global memory boards, and Environmental Control Module. These commands can be issued from the laptop service processor or, remotely, by the EMC Customer Support Center.

Also within the CCM is the top-level fabric for the Message Matrix communications system. This is a high-speed communications fabric within the Direct Matrix Architecture that allows for fast, reliable messaging between compute nodes.

Environmental Control Module

The ECM monitors and logs environmental events across all critical components and reports any operational problems. The ECM also allows the EMC Customer Product Lab to search for revision numbers and date codes on individual components. Critical components include Channel and Disk Directors, Global Memory Directors, Port Bypass Cards, power supplies, fans, and switches.

The ECM monitors each component's local voltages ensuring optimum power delivery. Fan speed is varied as required to maintain optimum operating temperature. The AC Power Main is checked for AC failures, transfer to auxiliary, DC failures, current-sharing between DC supplies, DC output voltage, specific notification of overvoltage conditions, the current from each DC supply, and voltage drops across major connectors. Ailing components can be detected and replaced before a failure occurs.

Non-Disruptive Operations

Enginuity is the intelligent operating environment for the Symmetrix DMX systems. Enginuity manages and ensures the optimal flow and integrity of information through the different hardware components of the Symmetrix DMX system. Enginuity manages all Symmetrix operations, from monitoring and optimizing internal data flow, to ensuring the fastest response to users' requests for information, to protecting and replicating data.

Enginuity provides the following services for Symmetrix DMX systems:

- Manages system resources to intelligently optimize performance across a wide range of I/O requirements
- Ensures system availability through advanced fault-monitoring, detection, and correction capabilities while providing concurrent maintenance and serviceability features
- Interrupts and prioritizes tasks from microprocessors to ensure that fencing off failed areas takes precedence over other operations
- Offers the foundation for specific software features available through EMC's disaster recovery, business continuity, and storage management software
- Provides functional services for both Symmetrix-based functionality and for a large suite of EMC storage application software
- Defines priority of each task including basic system maintenance, I/O and application processing including: EMC ControlCenter™, Symmetrix Remote Data Facility (SRDF), TimeFinder™, and EMC ControlCenter Symmetrix Optimizer
- Provides uniform access through industry-standard APIs for internal calls, and an external interface to allow integration with other software providers and ISVs

Enginuity upgrades and interim updates, performed at the customer site by the EMC Customer Engineer (CE), provide enhancements to performance algorithms, error recovery, reporting techniques, diagnostics, and code fixes.

Non-disruptive Enginuity upgrades from one version to the next, as well as interim updates (loads), are available for Symmetrix DMX, and take advantage of its multi-processing and redundant architecture. Release levels can also be loaded online without interruption to data availability.

During an online upgrade, the EMC Customer Engineer downloads the newest version of Enginuity to the Symmetrix DMX service processor. The new operating environment code loads into the EEPROM areas within the Channel and Disk Directors, but remains idle until it is made operational. The Symmetrix system does not require manual intervention to perform this function. All Channel and Disk Directors remain online to the host, thus maintaining uninterrupted application access. The Symmetrix DMX system will load the new version of Enginuity at selected windows of opportunity within each Director hardware resource until all Directors have been loaded.

Once the new version of Enginuity is loaded, internal processing is synchronized and the new code becomes operational. During an online load within a code family (or interim upgrade), the full version is loaded and consists of the same base code plus any additional patches.

The ability to perform non-disruptive Enginuity upgrades is critical to providing uncompromising levels of system availability and data access.

In addition to non-disruptive microcode upgrades, Enginuity also supports a wide range of non-disruptive operations allowing Symmetrix DMX to support mission-critical environments for applications that require uninterrupted access to information and uncompromising service levels. The following are examples of non-disruptive operations supported by Enginuity:

- Hardware upgrades, such as adding additional disk units, Channel Directors, and Global Memory Directors
- Configuration updates and modifications, such as reconfiguring or moving existing storage resources to support new host platforms
- Configuration updates and modification for TimeFinder and SRDF volumes, such as the ability to dynamically add or remove SRDF groups, or convert BCV capacity into standard devices

Symmetrix DMX Error Detection and Remote Support

Symmetrix DMX hardware is the most reliable in the industry.

However, all hardware is subject to occasional failures. The unique methods used by Symmetrix to proactively detect and prevent these failures set it apart from all other storage solutions in providing continuous data integrity and high availability. The Symmetrix system is designed with these end-to-end data integrity features:

- Error checking, correction, and data integrity protection
- Disk error correction and error verification
- Cache error correction and error verification
- Periodic system checks

Through the service processor, the Enginuity storage operating environment proactively monitors all end-to-end I/O operations for errors and faults. By tracking these errors during normal operation, Enginuity can recognize patterns of error activity and predict a potential hard failure. This proactive error tracking capability can fence off, or remove from service, a suspect component before a failure occurs.

The current dial-in support for the Symmetrix system uses the latest digital key exchange technology for strong authentication; layered application security; and a centralized support infrastructure that places calls through an encrypted tunnel between EMC Customer Support and the service processor.


Symmetrix component-level redundancy

All critical components are fully redundant, including Global Memory Directors, Internal Data Paths, Power Supplies, Battery Backups, and all Fibre Channel Back-End components.

Global Memory Director redundancy

The Symmetrix DMX Intelligent Global Memory Director technology (see Chapter Two) implements a fault-tolerant, highly reliable design.

Symmetrix DMX Global Memory Directors provide component-level redundancy where each component in the system is protected against failure. The redundancy is achieved through redundant components and via parity schemes. The following is an illustration detailing how a 64-bit word is protected via SNCDND (simple nibble correction double nibble detection) in a Symmetrix DMX Global Memory Director.



The connections from a region controller to global memory are monitored by the redundant region controller. Both region controllers are used to transfer data in the course of machine operations. At any given time, one region controller is performing the transfer while the other region controller is monitoring the data transfer to ensure that the protocol and signals are correctly asserted to the global memory components. In this operation, the two region controllers guarantee that, should they disagree, a fault will be detected.

Background diagnostics periodically check for faults. If a fault is detected, the Enginuity storage operating environment takes action to ensure that the data within that Global Memory Director is removed via the correctly operating region controller, and that it is “fenced off” and ready for replacement.

The global memory components themselves are also protected. Each component is just one in a RAID-like set. The Enginuity storage operating environment detects and corrects any possible component failures. Even if an individual component writes data to the wrong address, or does not perform the WRITE at all, the error is detected and corrected. In addition, when such errors are detected, Enginuity performs diagnostic tests to determine if the region has experienced a hard or soft failure. In the event of a hard failure, the global memory component is fenced off as before. In the event of a soft error, such as an alpha-particle-induced single-bit error, Enginuity logs the error, corrects the data, and continues to operate with that global memory. If the soft error rate exceeds a preset threshold, Enginuity fences the global memory for replacement.

Enginuity’s sophisticated fault-detection and correction capabilities, combined with the component-level redundancy characteristics of the architecture, provide superior data integrity.

Other aspects of machine design, including the extension of fault-recovery capabilities to every path in the system, further enhance data protection.

Memory chip redundancy

In Symmetrix DMX systems, each and every chip on the memory board is redundant. Memory chip redundancy eliminates any single point of failure on Global Memory boards, and produces higher system throughput through optimal utilization of global memory resources.

To achieve the goal of making every memory chip redundant, 8 bits of extra parity information are stored in addition to the usual 8-bits of parity information that go with a 64 bit long word. The result is 10 percent of extra memory capacity to create chip-level redundancy as compared to 50 percent waste in the case of mirrored cache boards.

Triple Modular Redundancy with Majority Voting

The Symmetrix DMX Global Memory Director deploys a state-of-the-art technique called Triple Modular Redundancy with Majority Voting (TMR-MV) across its two independent fault domains.

The Symmetrix implementation of TMR-MV uses a set of three identical modules created from components from three different vendors to eliminate any unstructured faults by ensuring that the output is that of the majority of modules.

Channel Director redundancy

Channel Director redundancy is provided by configuring multiple connections from the host servers (direct connect) or Fibre Channel switch (SAN connect) to the Symmetrix system. With SAN connectivity via Fibre Channel switches, each Symmetrix DMX port can support multiple host attachments, enabling storage consolidation across a large number of host platforms. The multiple connections are distributed across separate Channel Directors to ensure uninterrupted access in the event of a channel failure. A minimum of two connections per server or SAN is required to provide full redundancy.

EMC PowerPath

Channel failover functionality is required to automate failover and failback processes to avoid interruptions to data access. Without this functionality, a path failure due to a problem with the host bus adapter, Fibre Channel switch, fibre cable, or Channel Director would create the potential for the application to go down.

EMC PowerPath provides this functionality. It automatically detects when a path has failed and notifies the host that there is an inactive path. PowerPath then fails over the existing I/O request to another active path to maintain data access and application availability.

Once the failed path is fixed or repaired, PowerPath automatically detects that the path has become active, brings the path back into operation, and automatically starts sending I/O requests down the now active path.


PowerPath also provides intelligent load balancing to optimize performance and minimize bottlenecks. All this occurs transparently to the host so the application is not stopped and data is continuously available.

Internal control data path redundancy

Each Director has two paths to connect the Symmetrix DMX Message Matrix between each Director. The plug-in Communications Control Module provides independent redundant fault domains. In the event that a fault occurs within one of the domains the traffic is automatically routed by the hardware to the available domain and the system is alerted to the faulty CCM domain.

There are two CCMs per system (see diagram below). In a fully configured Symmetrix 2000, for example, 16 Directors are configured with 32 point-to-point serial connections. Sixteen connections are routed to the first CCM and the other sixteen redundant connections are routed to the second CCM.

The Message Engine (ME) component of the Message Matrix provides protection against rogue or run away nodes. Packet CRC and sequence numbers protect each message packet transmitted. Internal data paths between the Message Engine and Fabric Element are fully parity protected.



The Symmetrix DMX fully-redundant messaging matrix provides high-bandwidth and low-latency control communication to manage data movement throughout the Symmetrix system.

Power supply redundancy

Power supply redundancy in Symmetrix is achieved through a combination of redundancy techniques. In a typical application, an N+1 power supply system will be used. An N+1 system comprises multiple replications of small power supplies integrated together to satisfy the total power requirement of the system. If any single power supply fails, the remaining units can still maintain power to the system. N+1 also allows the failed power supply to be replaced without disruption to the system. Symmetrix DMX800 and DMX1000 systems implement an N+1 power supply.

In the Symmetrix DMX2000 and DMX3000, 2(N+1) redundancy is provided to eliminate these potential points of exposure.

2(N+1) power redundancy is based on dual N+1 power supplies for a higher level of fault tolerance and availability.

Single-phase or three-phase main and auxiliary AC power is provided to each of the two Power Line Input Modules (PLIMs).

The PLIMs feed main and auxiliary AC power to each of the DC power converters. Two independent Battery Backup Units support the DC power supplies, one for each zone. Critical boards receive power from each zone. Symmetrix systems have been designed to continue to operate in case of PLIM failure or replacement.

In addition to providing higher levels of availability, 2(N+1) redundancy is cost effective because the extra power supplied (beyond that needed with an N+1 implementation) is minimal.

Battery backup

The Battery Backup Unit (BBU) ensures data integrity during power outages. In the event of a power failure, the BBU allows the system to continue operation for several minutes. If power is not restored within that interval, the system is brought offline and powered down. The BBU delivers sufficient power to safely destage all cached data onto the disk drives. The BBU battery capacity is sufficient to ride out two consecutive outages. Approximately 12 hours of recharge time is required to fully recharge the BBU. The BBU also contains a self-diagnostic controller to ensure availability and correct operation by monitoring the battery condition and charger states, and by providing status updates to the Environmental Control Module.

Fibre Channel back-end redundancy

The Symmetrix DMX Architecture incorporates a two gigabit-per-second Fibre Channel back-end design to ensure high performance and full redundancy. The Fibre Channel back end provides redundant paths to the data stored on the disk drives to ensure non-stop access to information, even in the event of component failure and/or replacement. Within the Symmetrix back end, the Disk Directors, cable paths, disk ports and Port Bypass Cards are all redundant.

Arbitrated loop design

Symmetrix systems employ an arbitrated loop design that contains monitoring and control hardware and software to maximize the performance and availability of each loop. The loop is connected in a star-hub topology, with the hub ports gated with a bypass switch that allows individual loop segments to be dynamically inserted or removed. The loop initiator is a Symmetrix Fibre Channel Disk Director, which feeds data into and controls the hub. All of the monitoring and control logic is contained on the Port Bypass Card, which also contains the bypass/hub logic for the loop. To ensure the highest level of availability, the communications link to the monitoring and control functions is not carried by the Fibre Channel loop, but is implemented through a separate path. If the Fibre Channel loop is not operating for some reason, the Director can access the Port Bypass Card through the control bus and reconfigure the loop into a working state.

The Port Bypass Card is the heart of the Symmetrix dual-ported disk Fibre Channel loop and constantly monitors the loop for Fibre Channel layer 0/1 signal connection and protocol errors. If errors are detected, the controlling software program is notified, and if necessary the bypass circuits may be used to switch out the failing loop segment(s), thereby allowing the remaining devices on the loop to operate unaffected. Each disk enclosure contains dual Port Bypass Cards.

Each Port Bypass Card connects one of the two ports on the disk drives to a Disk Director. For each disk connection, a bypass circuit is provided that allows that disk connection to be shunted, either because 1] the drive is not installed, 2] it has just powered on or has been reset, or 3] it is malfunctioning.

In addition to the bypass circuits themselves, the card contains realtime monitors embedded within the loop that can determine if bad data is being generated on the loop. These monitors report status to the loop control software and can be useful in determining which parts of the loop might be subject to potential hardware failures.

The Port Bypass Card also has the ability to control power to the drives, and maintains basic configuration information. If the loop configuration is changed, the control software program is alerted so it can take appropriate action. The Port Bypass Card is located within the Disk Array Enclosure, which allows convenient replacement in case the isolation device itself fails. Each port of a disk is connected to a different Port Bypass Card, so access to the data is maintained even during service operations.

Note: Since the disk media is a single, shared resource, this level of redundancy will not protect the system from a disk media failure, a possibility that is best dealt with by RAID protection. Combining the two techniques, as is done in Symmetrix, yields the highest possible level of data access and reliability.

Symmetrix Component-Level Serviceability

Symmetrix DMX systems implement a modular design with a low parts count. This minimizes the number of failure points while improving serviceability by allowing non-disruptive component replacement should a failure occur.

Symmetrix DMX provides full component-level redundancy to ensure continuous and uninterrupted access to information. The design allows concurrent maintenance of all major components, including:

- Global Memory Director
- Channel Director
- Disk Director
- Disk devices
- Cooling fan modules
- Environmental Control Module
- Communication Control Module
- Power supplies
- Batteries
- Service processor

To improve serviceability, the Symmetrix DMX directors are mechanically keyed to prevent insertion to invalid slots. A “Disable Interface” feature ensures that at power up the link interfaces between any new Channel or Disk Director and its corresponding Global Memory Directors are disabled. Only after safeguards have been satisfied to ensure that a director has been properly replaced will a new director be activated and brought online.

Online Enginuity upgrades

Symmetrix Enginuity is the storage operating environment for Symmetrix DMX systems. Enginuity manages and ensures the optimal flow and integrity of information through the various hardware components of the Symmetrix system.

Enginuity manages all Symmetrix operations from monitoring and optimizing internal data flow, to ensuring the fastest response to users' requests for information, to protecting and replicating data.

Enginuity storage operating environment upgrades and interim updates, performed at the customer site by the EMC Customer Engineer (CE), provide enhancements to performance algorithms, error recovery and reporting techniques, diagnostics, and Enginuity fixes.

Online Enginuity upgrades from one version to the next and interim updates (loads) are available for Symmetrix DMX. Symmetrix systems take advantage of their multiprocessing and redundant architecture to allow for online microcode upgrades. Release levels can also be loaded online without interruption to data availability.

During an online Enginuity upgrade, the EMC Customer Engineer downloads the new version of Enginuity to the Symmetrix Service Processor. The new Enginuity loads into the EEPROM areas within the Channel and Disk Directors and remains idle until it is made operational. The Symmetrix system does not require manual intervention to perform this function. All Channel and Disk Directors remain online to the host, thus maintaining uninterrupted application access. The Symmetrix system will load the new version of Enginuity at selected windows of opportunity within each Director hardware resource until all Directors have been loaded.

Once the new version of Enginuity is loaded, internal processing is synchronized and the new code becomes operational. During an online Enginuity load within a code family (or interim upgrade), the full version is loaded. It consists of the same base code plus additional patches.

Symmetrix Data Protection Options

Symmetrix data protection options ensure uninterrupted access to data in the event of a disk failure. While disk drive technology has improved significantly over the last several years, the following disk protection options provide a higher level of data protection, recoverability, and availability. The options listed below can be purchased separately and implemented into the Symmetrix operation.

Data Protection	Description	Features
RAID 1 (Mirroring)	Data on one volume is protected by another identical mirrored volume to create a RAID 1 device.	• Highest performance—I/Os automatically distributed across both volumes • Highest availability—remaining volume available without rebuild • Efficient failed-drive recovery—single hot spare; drive copied upon replacement • Symmetrix Optimizer support
Parity RAID	Data on (3 or 7) volumes in Parity RAID are protected by a separate parity volume.	• High performance dependent upon volume and external striping • High availability—data from lost volume is regenerated from remaining members • Failed drive recovery requires N (3 or 7) hot spares; drive rebuilt upon replacement • No Symmetrix Optimizer support
RAID 5	Data and parity are striped across all (4 or 8) hypervolumes to create a RAID 5 device.	• High performance with automatic striping across hypervolumes • High availability—lost hypervolume regenerated from remaining members • Efficient failed-drive recovery—single hot spare; drive copied upon replacement • Symmetrix Optimizer support

Symmetrix DMX Mirrored and Parity RAID Options

Symmetrix Parity RAID technology is a combination of hardware and software functionality that improves data availability on drives by using a portion of the array to store redundancy information. This redundancy information, called parity, can be used to regenerate data if a disk drive becomes unavailable.

Within the same Symmetrix system, data can be protected through Mirroring, Parity RAID, or RAID 5.

Dynamic sparing can be added to any of these data protection options.

Disk mirroring concepts

Symmetrix mirroring provides the highest level of performance and availability for all mission-critical and business-critical applications.

Mirroring maintains a duplicate copy of a logical volume on two physical disk devices. The Symmetrix system maintains these copies internally by writing all modified data to both devices.

The mirroring operation is transparent to the host. The mirroring feature designates from two to four logical volumes residing on different physical devices as a mirrored volume, one volume being mirror-1 and the other volumes being mirror-2, mirror-3, and mirror-4. These mirrors can be local or remote. The host views the mirrored volumes as the same logical volume because each has the same unit address.

Parity RAID (3+1)

A Parity RAID (3+1) configuration consists of Parity Raid Ranks containing three data volumes and one parity volume. With this approach, effectively 75 percent of the total storage capacity of each Parity RAID group is available for storing data.

Parity RAID (7+1)

A Parity RAID (7+1) configuration consists of Parity Raid Ranks containing seven data volumes and one parity volume. With this approach, effectively 87.5 percent of the total storage capacity of each Parity RAID group is available for storing data.

Notes: Multiple Symmetrix RAID groups with one configuration can coexist on the same Symmetrix unit. Symmetrix RAID (3+1) groups and RAID (7+1) groups cannot be configured in the same Symmetrix system although each can be configured together with mirrored volumes. All logical volumes participating in a Symmetrix RAID group must have identical storage capacity. A Symmetrix system allows inter-mixing of different capacity physical disk devices within a single Symmetrix unit.

RAID 5 (3+1)

A RAID 5 (3+1) configuration consists of four Symmetrix devices with data and parity striped across each device. With this approach effectively 75% of the total capacity of a RAID 5 device is available for storing data.

RAID 5 (7+1)

A RAID 5 (7+1) configuration consists of eight Symmetrix devices with data and parity striped across each device. With this approach effectively 87.5% of the total capacity of a RAID 5 device is available for storing data.

Summary

Binding service level agreements commit IT organizations to deliver stipulated, measurable support metrics such as application performance, end-user response time, and system availability. Even in the absence of such SLAs, IT executives universally recognize that downtime can have disastrous ramifications in lost revenue, dissatisfied customers, and missed opportunities.

For over a decade, Symmetrix systems have been the gold standard for data integrity and availability in high-end storage.

The Symmetrix DMX Architecture raises the bar even higher, with enhanced availability features in every aspect of system design.

With key enhancements to a proven architecture, Symmetrix DMX systems are the logical choice for enterprises requiring only the most uncompromising levels of data and system availability for their high-end storage environments.

Chapter Five: Symmetrix DMX Functionality

In the course of the past decade, Symmetrix systems have proven to be the most capable storage system ever built. With the introduction of the Direct Matrix Architecture, Symmetrix functionality rises to a level never before seen in any high-end storage solution

Enginuity

At the core of the Symmetrix DMX Architecture, the world's most time-tested, proven storage operating environment

The Symmetrix DMX Architecture runs Enginuity, the most robust, proven, time-tested storage operating environment in the world.

More than a decade of accumulated algorithmic intelligence is hard at work in Symmetrix DMX systems from Day One.

In addition, Enginuity provides absolute application continuity with all existing Symmetrix environments, totally eliminating the need to retrain staff, rewrite applications, or do any new integration.

EMC's Enginuity storage operating environment manages the workload across Symmetrix platforms. It is a specialized storage operating environment that makes it possible for Symmetrix to serve as the foundation technology that ties together all elements of an IT infrastructure. The product of billions of dollars in R&D and engineering development spanning five Symmetrix generations, Enginuity contains more than three hundred and fifty patents.

An important benefit of this careful development is that all of the openness, reliability, availability, and serviceability features of Enginuity are present from the outset in each succeeding Symmetrix generation, including the Symmetrix DMX series systems.

Enginuity's unique value is built around four key concepts:

Foundation

Enginuity provides stability and continuity across technology generations, supporting powerful storage applications while insulating their functionality from technology changes.

Performance

Enginuity fully exploits the inherent performance potential of the Symmetrix architecture. Utilizing unique intelligent algorithms, it manages complex, dynamic environments to maximize system performance under any load.

Availability

Enginuity works incessantly and proactively to ensure data accessibility. It performs continuous integrity checking of data and hardware, detects problems, and implements controlled failover and escalation when they occur.

Openness

Enginuity sets the standard for broad integration and assured interoperability. Enginuity is an open environment using industry standard APIs. Hundreds of ISV applications can run on Symmetrix, even when new platforms are introduced.

Enginuity Foundation

The Enginuity storage operating environment sets EMC Symmetrix apart from all other storage technologies. Enginuity optimizes the capabilities of the architecture while enabling software applications to run across multiple generations of technology. Older and newer generations of Symmetrix will work together as one infrastructure, simplifying operations, saving time, lowering cost, and protecting investments.

Functions that distinguish the Enginuity storage operating environment from the simple microcode-based storage control designs of competitive systems are:

- Proactively orchestrating maximum performance, availability and data integrity
- Leveraging accumulated algorithmic intelligence
- Insulating applications from changing technologies
- Gracefully incorporating new technologies
- Allowing non-disruptive changes and upgrades to hardware and software
- Providing high levels of overall system security
- Executing all functions under pre-emptive and multi-tasking control

High-end storage systems today must support an increasing range of storage-specific applications, including local and remote replication; point-in-time copying; dynamic reporting and optimization; complex multi-fault recovery processes; and support for clustered application server environments.

Running these applications simultaneously will overwhelm any simple storage control design, as the number of events can dwarf routine I/O READ and WRITE processing.

Enginuity, on the other hand, as a storage-specific operating environment, is designed to support the event-driven multi-task processing that is characteristic of Symmetrix.

Together, Enginuity and Symmetrix can process simultaneous I/O requests from thousands of devices ranging from servers to switches, while taking escalating demands in stride.

Enginuity is also designed to maximize performance and operational efficiencies for people and processes as much as for technology and equipment. Enginuity ensures that users are never forced to relearn the entire system when underlying technology changes. Software application investments are protected because all processes and procedures can be fully leveraged over their entire lifecycles. Enginuity qualifies the advanced functionality of applications like SRDF on all new systems, for full compatibility among generations.

Security: Access Controls

Part of the Enginuity foundation is a sophisticated array of security measures that allow IT administrators to confidently manage access to data. Enginuity can be set to allow access to designated Symmetrix storage devices only by specifically designated hosts. Any other access attempt is rejected and recorded in security logs.

Enginuity approaches the challenge of strict data access control from multiple perspectives. First, Enginuity dictates physical points of access, i.e., which hosts, connected to which ports, can generate valid commands. Second, the system can be set to control who is allowed access and at what level of trust.

EMC ControlCenter software provides further access control by managing specific device pools within the array. Together, EMC ControlCenter and Enginuity define and manage administrative privileges at the finest level of granularity.

Enhanced System Audit and Investigation

Enginuity stores an Audit Log in the Symmetrix File System (SFS) that collects and presents a chronological list of host-initiated Symmetrix actions and activities. These include manual activities (such as removing/replacing a component), automatically initiated scripts, and API activities such as TimeFinder or Symmetrix Remote Data Facility (SRDF) routines.

Enginuity can provide a complete record of how and when a Symmetrix device is used. Also, the built-in system security of SymmAPI does not allow host applications to change audit logs. This means complete audit trail records are secure and available for any purpose.

Enginuity Performance: Dynamically Adjusting Performance Algorithms

Complex and constantly changing environments require continuous, intelligent, and aggressive management of all storage resources and functions. Enginuity uses advanced algorithms to maximize Symmetrix system performance. Working in concert with the Symmetrix architecture, Enginuity also makes special use of global memory to gain maximum performance and flexibility.

Storage systems use two types of memory: cache/global memory (fast memory) and disk (slow memory). Equivalent amounts of global memory are significantly more expensive than the slower disk, so utilizing both reduces cost. The goal of any storage system is to balance such use to:

- Maximize I/O from global memory
- Minimize the use of the disks for priority types of host workload
- Increase overall system performance
- Assure the ultimate in data safety

To do this, the system must 1] predict which pieces of information will be used farthest in the future, 2] store these on disk (the cheaper “slow memory”), and 3] manage user and application demands on the stored data to provide the fastest access to the greatest volume of information.

Enginuity optimizes the system so that required data is in the right place at the right time.

EMC Enginuity Performance Solutions

Enginuity apportions information and controls performance through dynamically adjusting performance algorithms. These are based on rigorous analysis of I/O usage patterns from thousands of installed Symmetrix systems as well as immediate conditions that apply to each individual system at any point in time.

By means of dynamically adjusting performance algorithms, Enginuity can ensure, with an extremely high degree of probability, that data will be readily available in global memory when needed.

Most storage vendors attempt to manage stored data through the classic LRU (Least Recently Used) algorithm. In this approach, the system looks at historical patterns of I/O access and tracks the time of access for each block. It then keeps in fast memory that information that has been most recently used. When the fast memory is nearly full, it flushes out the least recently used piece of information to the slow memory. This traditional approach makes controlling assumptions based on limited information, and severely compromises performance.

EMC performance algorithms can be grouped into three categories to address:

- Global Memory
- Disk
- Replication

Maximizing Global Memory Performance

The central challenge of managing performance in the faster global memory regions of memory is anticipating what data will be needed and making it available in global memory before users call for it. Other challenges include the need to ensure optimal distribution of global memory space to support application priorities, and to flag important data for fast retrieval even when it is needed only rarely.

Enginuity uses three types of algorithms to meet these challenges:

- Intelligent Adaptive Pre-fetch Algorithm
- Quality-of-Service (QoS) Algorithms
- Permacache

Intelligent Adaptive Pre-fetch Algorithm

When users access information in predictable ways, storage performance can be improved by “pre-fetching” information to global memory from disk before the host actually requests it.

Pre-fetching can dramatically help performance because response time improves by a factor of ten if the data resides in global memory when the host needs it. It also improves disk utilization by transferring large portions of data each time, virtually eliminating seek and latency delays.

The EMC Enginuity Solution

Enginuity takes a sophisticated approach to selecting which information to fetch in advance. The starting point is a statistical foundation that automatically adjusts to the workload by constantly monitoring and evaluating the success of its decisions. By incorporating this added intelligence into pre-fetch algorithms, Enginuity pre-fetches the data that is needed, on time, without affecting the response times of other I/O events.

Quality-of-Service (QoS) Algorithms

With multiple demands placed on a storage platform by a range of customer applications, a user will invariably prioritize those applications according to their importance to the business. But if the storage platform can't distinguish among the applications, those priorities are meaningless in terms of IT performance. In the absence of intelligent prioritization, for example, a busy retail website during the holiday shopping rush will get the same priority as back-office processing.

The EMC Enginuity Solution

QoS is enabled by allocating a portion of global memory for a subset of logical volumes. This enables users to control the hit ratio for a certain application (and the associated group of logical volumes) regardless of other applications running on the same system. QoS also enables users to lend a portion of global memory to other applications. QoS is implemented by a mechanism called multiple priority queue that allows application performance to be adjusted to match business needs.

Permacache

With conventional pre-fetch approaches, data that is rarely accessed resides by default in the slower-performing disk.

However, some of this data, called High Priority, Rarely Accessed data (HPRA) requires special treatment because, although rarely accessed, it must be immediately available when needed.

The EMC Enginuity Solution

Enginuity enables users to dynamically assign (and de-assign) a portion of global memory slots as Permacache for use with HPRA data.

Maximizing Disk Performance

The mechanical characteristics and limitations of disk drives require a special set of algorithms to obtain optimal performance and provide a satisfactory complement to the faster global memory.

Techniques used by Enginuity to maximize disk performance include:

- Dynamic Mirror Service Policy
- Fast Write Algorithm
- Rotational Position Ordering

Dynamic Mirror Service Policy (DMSP)

Mechanical movement of the spinning disks requires a wait time while the disks rotate and align themselves to the magnetic heads inside the disk drives. The radial movement of the heads adds more wait time as they seek out a particular cylinder. Mirroring, using two disks, can double wait time and compromise READ performance. When an operating system tries to resolve this problem by assigning one mirrored disk to serve as a source for READ operations, the spindle associated with the secondary disk is idle and underutilized.

The EMC Enginuity Solution

For mirroring, Enginuity collects statistics on spindle rotation and disk head seek movements over a predefined interval (usually five minutes). Enginuity uses this information to decide between “split” versus “interleave”¹ reading patterns. This informed decision-making optimizes READ performance while using mirrored disks.

Rotational Position Ordering (RPO)

A “first-in/first-out” system of disk access produces unacceptable seek and latency times associated with mechanically rotating spindles. Traditional storage system microcode does not improve upon the sub-optimal seek and latency times and degraded I/O performance that comes from servicing I/O in the order of their arrival.

The EMC Enginuity Solution

Whenever multiple I/O requests are queued on the disk, the drive microcode optimizes the order in which I/Os are executed by the disk.

Disk Pre-fetch Control allows a Symmetrix system to pre-fetch more aggressively from disk to global memory. Intelligent algorithms, working between disks and global memory, pre-fetch data only when there is a high probability that the host will read the data in case of a sequential I/O.

Enginuity Availability

Enginuity is a multi-tasking, event-driven storage operating environment. Throughout the system, it continuously monitors dozens of elements that check data integrity and performance. It watches the results from checks on data integrity in global memory, disks, and across connections. It monitors hardware components for voltages, stability and trends. And it does all of this in the background while simultaneously continuing to prioritize and perform the work of data storage and movement at ultimate performance. Enginuity optimizes availability by:

- **Continuous integrity testing and error detection**

Through continuous self-monitoring and self-diagnosis of the entire Symmetrix system, Enginuity ensures that information is always available. It manages and tests all hardware components continuously, and automatically makes adjustments to maintain optimal operational hardware efficiencies.

¹ SPLIT—A system of logically splitting a disk into two halves, the inner cylinders and the outer cylinders. The I/O corresponding to the inner part of the mirrored disk is serviced by disk 1 (primary disk) and the I/O corresponding to the outer part of the mirrored disks is serviced by the disk 2 (secondary).

INTERLEAVE—A system of logically reading alternate tracks from mirrored disks. E.g., Enginuity reads even numbered tracks from disk 1 (primary disk) and odd numbered tracks from disk 2 (secondary disk).

- **Continuous trending analysis**

The Enginuity storage operating environment is finely tuned to relevant events in the Symmetrix system. It is extremely intelligent in interpreting which events may be significant. It constantly probes for potential problems, detects trends, predicts faults, and takes appropriate and autonomous corrective action.

- **Failure containment and recovery**

In severe circumstances, a fault may require Enginuity to intelligently fail over to redundant components. It then initiates a call home for service or replacement. It does this transparently to users and with no adverse impact to the enterprise. When repairs are complete, Enginuity seamlessly brings all components back to normal operation.

- **Non-disruptive upgrades**

In a dynamic business environment, there is an inherent need to grow and re-configure hardware, re-configure or add software or software patches, and do it without disrupting business operations. Applications must continue to function with the smallest possible degradation of performance for the least amount of time. With each new version, EMC has improved Enginuity's capacity to carry out more upgrades online without disturbing business operations.

Continuous integrity testing and error detection

Enginuity monitors and controls Symmetrix hardware to detect and resolve potential availability issues with no impact to system operation. Users are typically unaware that a problem ever existed.

For example, continuous testing includes checking for faulty disk sectors, comparing signals from identical redundant modules, and monitoring watchdog timers and voltage levels. Enginuity protects all the information flowing through the system via multiple techniques.

At every critical point in the system, Symmetrix deploys identical modules across all paths. So, for instance, Symmetrix Global Memory Directors are designed with two independent fault domains, each with its separate system clock, logic and redundant power supplies. Enginuity then uses signal comparison techniques to constantly monitor hardware and ensure data accuracy at every point.

Enginuity monitors watchdog timers to detect errors. Symmetrix loads these with a time value to complete a given operation. As this occurs, the watchdog timer is reset for the next operation. If the operation does not complete within the designated time, a fault is detected and Enginuity takes appropriate fault recovery measures.

Parity Check

Multiple parity bits are associated with every I/O flowing through the Symmetrix. The parity bits are checked for validity at every stage to ensure the information is interchanged.

Checksums and Cyclical Redundancy Checks

A check sum and Cyclical Redundancy Check (CRC) are generated and appended to every block of information. These are recalculated and checked for accuracy at each stage the information block is moved.

Continuous trending analysis

Advanced fault-tolerant techniques implemented by Enginuity work continuously to assure that storage operates to specification in spite of existing faults. Faults can be permanent, transient, or intermittent. They can also be structured or unstructured. Faults can be detectable, correctable, or recoverable. Faults may result in error conditions that are detected and corrected or recovered. In the case of Symmetrix with Enginuity, constant monitoring tracks minor faults. Trends are identified and future, more severe faults can be predicted and prevented. They will not impact system performance because they will not occur.

Failure containment and recovery

Proactive Cache and Disk Scrubbing

Enginuity periodically reads and rewrites all areas of global memory to detect any errors. This cache scrubbing technique maintains a record of errors for each memory segment. If the predetermined error threshold is reached in a certain segment, Enginuity directs the service processor to generate a call-home for immediate attention.

Constant cache scrubbing is employed to detect and correct occasional small “soft” errors, dramatically reducing the potential for larger “hard” errors. If a permanent error is confirmed, it is immediately contained by removal from service. Where feasible, the segment’s contents are moved to another area in global memory. Enginuity then continues to monitor conditions and, if an unacceptable level of errors is detected, it directs the service processor to call home for a non-disruptive memory replacement. As soon as the faulty part is replaced, Enginuity sees it and restores normal activity.

A similar technique is deployed for all data blocks on Symmetrix drives. During a production lull, Enginuity reads and rewrites all blocks on all tracks on all physical disks and then checks for errors or inconsistencies. Bad sectors may be walled off or, if significant disk errors are reported, the service processor generates a call home. As soon as the faulty part is replaced, Enginuity sees it and restores normal activity.

Checking the checkers

Beyond the comprehensive integrity checking functions described above, Enginuity also verifies that the checkers themselves are operating correctly. To accomplish this, Enginuity introduces benign errors. If the benign errors are checked properly, the checkers are considered to be working properly. Otherwise, the checker itself is marked faulty, a flag is raised to initiate a call home, and the component is replaced.

Instruction monitoring

The design philosophy of both Symmetrix and Enginuity is to detect and correct errors at every stage. To this end, the address component of every command/instruction is checked for validity to ensure that the operation is being undertaken at the correct address location in global memory. If an error is detected, the instruction is discarded and a re-try message is sent back to the host.

Non-Disruptive Upgrades (NDUs)

Enginuity 5670+ provides non-disruptive upgrades from Enginuity 5670 or 5669 across code families. As new features are incorporated into future Enginuity versions, they will be engineered to enable non-disruptive upgrades to succeeding versions.

Change Management

All Enginuity upgrade and configuration changes are enabled only in the context of a careful change management process that keeps application environments functioning and avoids error conditions. In some cases during an upgrade, the application may function with slightly degraded (up to 15 percent) service for a short period of time.

Enginuity's Open Environment

An enterprise storage array like Symmetrix must communicate with many different network devices, hosts, and network components. It must deal with clustering, failover, storage applications, and multiple versions of many different operating systems.

Achieving interoperability can be a major challenge. As the complexity of IT environments increases, solving interoperability problems becomes even more difficult.

Many factors come into play here, any one of which can spell disaster if not properly accounted for. The industry-accepted Fibre Channel protocol, for example, allows for significant flexibility in actual implementation. As a result, vendors implement the code in their own ways, which leads to inconsistencies in how devices interact, with potential interoperability issues across a storage network.

EMC's interoperability program prevents this kind of technology conflict. Dozens of operating systems, in multiple versions, with hundreds of different servers connected to thousands of network devices, are tested for operational integrity when attached to a storage network.

Open Standards

Enginuity integrates standards that incorporate both EMC and industry technology. Combined with EMC open application programming interfaces (APIs), this helps partners and customers to more rapidly develop standards-compliant, open storage application solutions.

EMC has made APIs available to partners since 1994. These APIs and, where appropriate, command line interfaces (CLIs), provide access to the sophisticated functionality of Enginuity and Symmetrix. They make it possible to develop extended storage applications and they can also interface with the most widely used enterprise management frameworks. The result has been hundreds of separate development efforts leading to valuable third-party storage applications. In addition, Enginuity's open APIs protect investments in these applications as the interfaces carry forward with successive versions.

EMC is committed to incorporating the emerging SMI specification so that applications that manage SMI-compliant devices can be quickly and securely developed.

As a result of EMC's commitment to open standards, powerful software can orchestrate all hardware, onboard functionality, and application workloads concurrently while maintaining the highest levels of end-user responsiveness and system availability.

Networked Storage


Networked storage from EMC leverages industry-leading Symmetrix functionality across the entire IT infrastructure, without geographical boundaries. It unifies storage networking technologies, storage platforms, software, and services so the enterprise can efficiently and cost-effectively manage, protect, and share information.

Among the benefits of networked storage are the cost savings associated with consolidating servers and storage. The more servers and storage that are eliminated, the greater the benefit. With the Symmetrix DMX Architecture, the storage for thousands of servers can be consolidated onto a single Symmetrix DMX storage array with dozens of terabytes of capacity while maintaining high-end service levels for performance, availability, and functionality.

Although networked storage is a simple concept to understand, it is very difficult to implement. In order to deliver a complete networked storage solution, EMC does far more than simply put a Fibre Channel switch between storage and servers, or attach a storage array to an IP network. EMC has made a commitment to integrate best-of-breed SAN and NAS technologies and ensure that they all interoperate together seamlessly.

Beyond the interoperability of hardware from multiple vendors, EMC also has the software capabilities to control and manage all of these devices as the environment grows and requirements change.

EMC's term for the integration of SAN and NAS technologies with the software to manage it all and the services to keep the environment up at all times is Automated Networked Storage. EMC is the only storage vendor in the industry 100 percent committed to this concept.



EMC Automated Networked Storage creates an integrated, managed storage network that delivers three benefits: consolidation, control, and continuity.

Interoperability

As enterprises increasingly deploy networked storage as the foundation of information-centric environments, the size and scope of the interoperability challenge grows.

Interoperability problems reach far beyond hardware connectivity. Users can invest countless hours and valuable resources to ensure that all components work together, from switches, routers, and host bus adapters (HBAs), to operating systems and applications.

EMC has invested nearly ten years and \$2 billion in resources to qualify and test the widest possible range of components for interoperability across networked storage environments.

EMC performs the industry's broadest, deepest, and most exhaustive regimen of multi-vendor interoperability testing. Employing world-class specialists, facilities and methods, EMC's E-Lab™-tested interoperability program spans the multi-vendor information storage infrastructure, which includes:

- Operating systems, protocols, and file systems
- Switches, host bus adapters, and other connectivity devices
- Storage devices, including disk- and tape-based systems
- Databases, application software, backup and recovery solutions, and third-party storage management software

To date, EMC has tested and qualified interoperability between its storage systems and nearly 400 server models, 40 operating systems, 81 storage software products, 145 networking elements and 1,200 devices ranging from HBAs and drivers to switches and tape subsystems. And this testing continues today. This comprehensive approach ensures that multi-vendor interoperability can be sustained even in the face of new software revisions, hardware upgrades, and operating system changes.

Networked Storage Evolution: Driven by Consolidation

The driver for networked storage has always been consolidation.

The ability to consolidate many servers onto a single storage array through a storage area network improves utilization and simplifies a number of functions, most notably business continuity.

Storage area networking (SAN) redefined the storage infrastructure by allowing more servers to attach to multiple storage arrays using Fibre Channel technology.

IT organizations found that consolidation brought significant benefits: better resource utilization with improved data availability, protection, management, and sharing.

Like SAN, networked attached storage (NAS) also became a buzz phrase in the industry, manifesting at first as a stand-alone storage device designed to share data over an IP network.

Today, EMC's implementation of networked storage delivers an integrated SAN and NAS solution leveraging the unique benefits of each technology.

EMC takes the model even further with Automated Networked Storage, an intelligent infrastructure that proactively manages the environment to minimize duplication and simplify administration.

To fully appreciate the benefits of an integrated SAN and NAS solution, it helps to look at the unique advantages that each brings to the table.

Storage Area Networks

Fibre channel-based storage area networks are ideal for applications that require high performance and throughput.

Fibre Channel SANs support block-based I/O between the host and storage, which is typically a requirement of databases and OLTP applications demanding high I/O performance.

Fibre Channel, today running at 2GB/s, offers a throughput advantage over previous single-channel SCSI direct-connect technology.

Fibre Channel can also extend the distance beyond a typical SCSI connection by up to 100km using technologies such as broadband, Dense Wave Division Multiplexing (DWDM), and Fibre Channel over IP.

With its high-speed Fibre Channel connectivity, SAN creates a dedicated storage network specifically designed to share physical infrastructure. Among the many benefits of SAN, the most dramatic are cost reduction through consolidation and enhanced backup and recovery.

Cost reduction through consolidation

The traditional relationship between a server and storage is one to one. This is the direct attached storage (DAS) model. When storage requirements grow beyond the capacity of the server, a second server is added.

SAN allows for the decoupling of this one-to-one relationship. Instead, storage can be allocated to servers as needed by leveraging a common, shared domain. The application environment can grow independently of the storage environment, and vice versa.

The ability to share back-end storage infrastructure across multiple platforms increases utilization. The total amount of storage bought, but unused, goes down.

Moreover, consolidation creates an information-centric infrastructure, where similar processes can be applied across the environment regardless of the type of application on the server.

Backup and recovery

In a DAS environment, backup and recovery is a complex process that becomes extremely costly over time. Typically, many tape libraries are deployed.

When a tape library is attached to a SAN, however, it becomes a shared resource that can be used across multiple applications and servers, simplifying backup processes and reducing their cost.

SAN in a Symmetrix DMX environment

The Symmetrix DMX Architecture, combined with EMC Connectrix Fibre Channel directors and switches, delivers unparalleled scalability and flexibility for the storage area network.

Multiple Connectrix Directors and switches can be mixed and matched in any combination to build a SAN precisely tailored to the requirements of the environment. Many servers can be connected to a few Symmetrix systems, or a few servers can be connected to many Symmetrix systems. The payoff is seamless, end-to-end interoperability across the entire network.

Networked Attached Storage

Network attached storage is ideal for sharing information such as Windows or UNIX files. Most companies have multiple sharing environments such as home directories, Web content, software development and design. A key requirement for NAS is to provide a single storage solution that integrates the entire environment, no matter how diverse its needs might be.

The most common way to share files is with a general-purpose server connected to an Ethernet network. These are inexpensive and easy to install. Difficulties arise, however, when the environment has to expand and more servers are added. As the number of general purpose servers increases from two to tens to hundreds, their low acquisition cost is more than negated by the rising expense of managing them. In addition, administrators find themselves spending more and more time at management tasks: licensing and maintaining each individual server, clustering software to provide high availability, and attempting to implement backup processes without impacting network performance. The answer to these challenges is network attached storage from EMC.

The Celerra family of NAS products provides a broad range of consolidation options based on Symmetrix and CLARiiON platforms.

Celerra systems are designed for consolidation and optimized for moving data. They offer enterprise-class NAS within the data center, regional office, or workgroup.

The Celerra File Server delivers high-performance information sharing among heterogeneous networked servers and end-user clients across any distance. In addition to superb performance, it brings unprecedented levels of availability, data protection, ease of management, and scalability to network file storage.

Integrated SAN and NAS

Networked storage integrates SAN and NAS in a way that leverages the consolidation benefits of both in a single solution.

Simply connecting the two technologies isn't enough. The point is to combine them into a unified environment with common processes and procedures, the same management tools, and the same back-end storage functionality.

The combination of a SAN using Symmetrix, and NAS leveraging Celerra, delivers unprecedented levels of scalability and integration in networked storage. Among the benefits of EMC's SAN and NAS integration strategy are:

- The highest levels of file server consolidation in the industry, eliminating the need to grow large file serving environments that are costly and cumbersome to manage
- The unique capability to allocate storage across SAN and NAS environments, eliminating the need to have two storage domains for different application types
- The ability to leverage the same backup infrastructure on the SAN for NAS file systems
- The ability to leverage the same remote replication process used for SAN-based applications for NAS file systems as well
- The ability to flexibly allocate storage for file shares, for e-learning, software development, CAD/CAM, Web content, databases, OLTP—all from a single infrastructure
- The ability to scale flexibly by allocating from the SAN or adding NAS components to the SAN

Networked storage from EMC combines the unique advantages of both SAN and NAS into a centrally managed, single infrastructure offering the lowest TCO in the industry.

EMC's networked storage strategy is future-ready. The Symmetrix DMX Architecture is designed to provide a broad range of connectivity options, and to incorporate new ones as they are introduced and accepted into the marketplace. Today, EMC Symmetrix supports Fibre Channel SAN, FICON, or ESCON for the mainframe, or NAS for information sharing. As technologies change, the Symmetrix DMX Architecture, with the intelligence of the Enginuity storage operating environment will facilitate the seamless integration of future technologies.

EMC Business Continuance Software


The TimeFinder and SRDF families of local and remote replication solutions deliver the most comprehensive and robust suite of replication solutions available in the marketplace, providing high performance, a wide range of deployment options, and an industry-proven architecture. The TimeFinder and SRDF families of remote replication solutions enables organizations to balance performance, availability, functionality, and economic requirements to achieve required service levels for local and remote disaster recovery and business continuity.

TimeFinder/Mirror is software for the Symmetrix that facilitates creation and maintenance of local multiple mirror images of production data without disruption to the primary application. The images are referred to as business continuity volumes and are widely used for backups, data warehouse loading, decision-support applications, application development, and rapid data restoration.

TimeFinder/Mirror can provide near instant recovery from software failures, operator errors, or hacker attack. TimeFinder BCVs are made using Symmetrix resources and require no host, server, or channel resources. They can be implemented in an efficient, repeatable, affordable, and consistent manner regardless of the architecture, file system, or database manager of the source data.

TimeFinder can also integrate backup, management, and data extraction tools from many EMC partners via Symmetrix APIs.


TimeFinder/Mirror business continuance volumes are full copies of production information that can be used for parallel processing activities such as backups, decision support, reporting, and/or application testing and development.



Symmetrix Remote Data Facility/Synchronous (SRDF/S) and SRDF/Asynchronous (SRDF/A) provide mirroring between Symmetrix arrays, which may be a few meters or thousands of miles apart. By maintaining data at a second site, SRDF/S and SRDF/A facilities provide rapid business resumption in case of disaster or data center relocation. Traditional data recovery tasks, such as tape restore and database recovery, which are labor-intensive, error-prone, and extremely time-consuming, are eliminated.

SRDF/S and SRDF/A are host (server), operating system, file system, and database independent. They represent the only solutions of their kind that work simultaneously with mainframe, UNIX, Windows, and Celerra systems.

SRDF and SRDF/A enable you to achieve continuous business operations via remote failover and failback of your environment in the event of a planned or unplanned site outage.



TimeFinder/Snap for Symmetrix creates space-saving snapshot images instead of full-volume copies on Symmetrix DMX systems.

Up to sixteen snapshots can be created from a single source volume, with each snapshot consuming minimal space. TimeFinder/Snap supports the ability to restore the saved point-in-time image to the source volume, to a TimeFinder/Mirror business continuance volume (BCV) of the source volume, or to a separate device.

In addition to facilitating recovery and restart operations, TimeFinder/Snap also expedites application development by providing developers with multiple, consistent views of the same data.

Alternative solutions for long-distance replication

Some IT organizations are reluctant to implement SRDF/S over long distances because extended distance raises two concerns:

- 1] Synchronous remote mirrors will inevitably cause a performance problem for any reasonably active update application.
- 2] Even if response time is not an issue, the cost of a long-distance, high-bandwidth communications facility may be too great to make the implementation feasible.

Among the long-distance replication solutions available to Symmetrix users are Symmetrix Remote Data Facility/Asynchronous (SRDF/A) and Symmetrix Remote Data Facility/Automated Replication (SRDF/AR).

Symmetrix Remote Data Facility/Asynchronous (SRDF/A), designed for both open systems and mainframe environments, is intended for use over greater distances than are feasible with synchronous replication technologies. It delivers a consistent and restartable remote copy of production data at all times, over any distance, with no host application impact, while minimizing bandwidth requirements. SRDF/A is able to accomplish this because of EMC's patented Delta Set architecture.

Delta Sets are global-memory-resident collections of WRITES that have occurred within a specific period of time. Bandwidth is conserved because Delta Sets allow data to be re-written, and only the final set of updates is sent over the communications link to the remote site.

In addition, SRDF/A allows the IT administrator to size communications links to the average WRITE workload rather than the peak workload. Because Delta Sets are collecting the new WRITES in global memory, SRDF/A is able to ride through periods of increased WRITE activity more efficiently.

Symmetrix Remote Data Facility/Automated Replication (SRDF/AR) option

Symmetrix, combined with its industry-leading SRDF/S and TimeFinder/Mirror software along with optional SRDF/AR software, allows users to implement automated regional failover combined with long distance replication with no data exposure while:

- Masking the effects of distance with respect to performance
- Minimizing communication costs
- Facilitating rapid restart of applications at the recovery site with no data loss

SRDF/AR can be configured in a single-hop or multi-hop configuration. Single-hop uses two Symmetrix systems with SRDF/S and SRDF/AR along with TimeFinder/Mirror and automates the cycling of TimeFinder/Mirror BCVs on the Source side and sends the scheduled updates to the Target system.

Multi-hop is implemented by placing by an additional Symmetrix array, with TimeFinder/Mirror, between the source and target Symmetrix arrays. (Servers are not required at this intermediate site.) In this way, long-distance, high-capacity bandwidth is replaced with short-distance, low-capacity, inexpensive bandwidth.

SRDF/AR Multi-hop works as follows:

- SRDF/S maintains data synchronization between the production site and the intermediate site. A TimeFinder/Mirror BCV is established to the target volume in the intermediate Symmetrix.
- At regular intervals, the BCV is split from its source, and in turn becomes a source volume, synchronizing with a target volume in the recovery site. This synchronization is incremental, so the bandwidth requirements between the intermediate site and the recovery site are far less than would be required in a non-buffered implementation. Once synchronized, SRDF/S or SRDF/DM is suspended between the intermediate and recovery sites.
- The BCV is reestablished to its source volume, and is incrementally resynchronized with updates that occurred while the BCV was synchronizing to the recovery site.

The entire process repeats in an automated fashion by SRDF/AR on an established schedule. The SRDF/AR capability makes this implementation very easy.

Symmetrix Remote Data Facility/Data Mobility (SRDF/DM) uses an Adaptive Copy mode to provide data mobility and/or migration between two or more Symmetrix systems.

SRDF/DM is ideal for data exchanges and content distribution. Adaptive Copy mode enables applications using that volume to avoid propagation delays while data is transferred to the remote site.

SRDF/DM supports all Symmetrix systems and all microcode levels that support SRDF, and can be used for local or remote transfers.

Note: Unlike full-function SRDF/S, SRDF/DM is not intended for disaster recovery.

SRDF/Consistency Group (SRDF/CG) is a software option designed to ensure the consistency of data participating in an SRDF family session.

Most applications, and in particular database management systems, have dependent WRITE logic embedded in them to ensure data integrity if a host processor, software, or storage subsystem fails. An example is a database update. When a DBMS updates a database, it first writes to the disk containing the log, then to the actual database dataset, and finally to the log volume to indicate that the update was made. The three WRITE I/Os (log, database, log again) are related, and each I/O is not issued until the prior I/O has been successfully completed.

In a remote disk copy environment, data consistency cannot be ensured if one of these I/Os was remotely mirrored, but its predecessor was not. This could occur, for example, in a rolling disaster, where a communication loss affects only a subset of the disk controllers that are performing the remote copy function.

EMC's SRDF/Consistency Group option prevents this from happening, by intercepting any I/O to a volume that cannot communicate to its remote mirror. It then suspends the remote mirroring for all volumes defined in the consistency group before completing the intercepted I/O and returning control to the application. By preventing independent I/O from being issued by the application, a consistency group implementation ensures the integrity and consistency of the data at the remote site.

The Symmetrix Direct Matrix Architecture software performance

The Symmetrix DMX Architecture dramatically enhances the performance of EMC business continuance software.

TimeFinder/Mirror activities are significantly faster. Users can establish more BCVs, split and resynchronize more often, and provide better service and more current data for TimeFinder/Mirror applications. Also, TimeFinder/Snap's space-saving snapshots on DMX systems enable users to reduce recovery point objectives while utilizing less disk capacity than ever before.

SRDF family users benefit in several ways:

- The added flexibility of Symmetrix DMX configurations makes it possible to have more SRDF-based Remote Adapters without compromising front-end connectivity. This, in turn, makes it possible to have more data and more parallelism of data transfers, so more applications can be affordably protected by SRDF or SRDF/A.
- With the power of the Symmetrix DMX Architecture, the part of SRDF/S and SRDF/A operation that take place within the array are greatly accelerated. For activity with relatively short communication paths, this makes more applications eligible for remote replication protection as overall response times are reduced.
- SRDF/AR in a multi-hop configuration, or periodic refreshes of data from TimeFinder/Mirror BCVs, can be expedited. Because TimeFinder/Mirror reestablish and split times are so dramatically fast, cycle times for these refreshes also can be shortened with significant gains in performance, granularity, and throughput.

Open Storage Management

A Symmetrix-based storage environment is superbly functional, thanks to the Symmetrix DMX Architecture, Enginuity storage operating environment, Symmetrix software, and EMC's networked storage connectivity.

However, the real world environment is seldom Symmetrix-only. It is more likely to be multi-vendor and diverse, and to have mid-tier as well as high-end storage requirements.

A dispersed, diverse, multi-vendor environment can be notoriously hard to manage, and, functionally, may appear to beleaguered administrators to be far less than the sum of its parts.

EMC's Open Storage Management software strategy was specifically devised to extend Symmetrix functionality and ease of management throughout a diverse, multi-vendor environment. It includes three complementary initiatives:

- **Intelligent Supervision**
- **Information Safety**
- **Infrastructure Services**

Intelligent Supervision

EMC's ControlCenter family of products automates the monitoring, reporting, and control of a networked storage infrastructure, with an emphasis on managing resources, networks, and devices.

Resource Management involves 1] monitoring, reporting, and controlling performance utilization; and 2] allocating resources across the environment, from user applications to network storage.

Effective resource management is key to driving service levels up while driving costs down.

From an *asset perspective*, it is essential to know how much storage is available, who's using it, and when more will be required.

The problem with traditional, non-automated methods of resource management is that they are labor-intensive and slow. By the time the required data is accumulated and analyzed, it may no longer be accurate or relevant. A far better approach is to automate resource management tasks, saving time, achieving greater efficiency, and freeing IT staff to focus on initiatives that drive the business forward.

From a *performance perspective*, the primary areas of concern are performance tuning, capacity planning, and problem isolation. Storage administrators need to correlate host performance, SAN and storage device performance information in a holistic rather than piecemeal fashion.

EMC's resource management software automates the process of discovering and reporting; monitors resources in real time; and generates reports for planning or charge-back purposes.

EMC resource management products include:

- *Automated Resource Manager*: for simplified, automated provisioning and alert management, and consolidation of storage resource operations
- *StorageScope*: for enterprise-wide storage utilization reporting and planning
- *Workload Analyzer*: for server, SAN, and storage performance analysis
- *Database Tuner*: for performance reporting on Oracle, UDB, and Microsoft SQL databases

Given the complexity of SAN environments, storage area network management has been a time-consuming task requiring many tools and much expertise. And the stakes are high. When networks aren't configured properly, bottlenecks are created between application servers and storage.

EMC's answer is SAN Manager, which delivers consolidated, multi-vendor SAN management including LUN masking, fabric zoning, and auto-pathing.

Device Management

In medium to large environments, there may be dozens of different tools that are being used for tasks such as discovery and monitoring, and for configuring storage devices and NAS devices.

All of these disparate tools create an extremely complex environment. Each tool has its own specific view that does not take other devices and management products into account. And each has its own look and feel, so that training and processes for management cannot easily span the entire environment.

The result of all this complexity is predictable: an inefficient operation that does not scale well.

EMC provides the ability to consolidate device management.

Managing disparate tools through a single interface simplifies the storage administrator's task and reduces the need for training by giving all management applications a common look and feel.

EMC device management products include:

- *Navisphere*: for monitoring, provisioning, and reporting on CLARiiON storage arrays
- *Common Array Manager*: for monitoring and reporting on third-party storage arrays
- *Symmetrix Manager*: for monitoring, provisioning, and reporting on Symmetrix storage platforms
- *Symmetrix Optimizer*: for automatic optimization of volume performance on Symmetrix

Information Safety

Information Safety in a heterogeneous, multi-vendor environment requires efficient, centralized management of three key areas: recovery, replication, and information lifecycle management.

Recovery Management means being able to withstand unexpected events and minimize the time needed to restore business operations.

Replication Management means enhancing information protection and increasing information utilization by enabling parallel business operations.

The software products in these two groups are used for information recovery or information replication, depending upon how they are deployed. These products include:

- *EMC Data Manager (EDM)*: for automating enterprise-wide backup and recovery for open systems
- *Replication Manager*: for automating the management and scheduling of TimeFinder, SnapView, and third-party data replication products
- *Symmetrix Data Mobility Manager*: for automating the scheduling and use of remote and local information replication
- *SRDF/TimeFinder Manager*: for monitoring and controlling local and remote Symmetrix mirrors through the ControlCenter console

EMC's policy-based **Information Lifecycle Management** products are designed to meet current and changing business requirements by moving information seamlessly across networked storage.

Information can be stored on the most appropriate storage platform for its current use, and can be moved as it ages. This lifecycle-based movement of information is automatic and based on business policies and is transparent to end-user applications.

Infrastructure Services

The Infrastructure Services initiative provides for the logical presentation and seamless mobility of information.

These products further hide the details of the storage infrastructure and provide policy-driven services that:

- provide logical presentation (abstraction) of physical devices
- automatically optimize the use of different physical storage elements and topologies
- deliver these services seamlessly to the application environment

Infrastructure services are focused on three areas: access optimization, data presentation, and data mobility.

Access Optimization is achieved with intelligent multi-pathing tools providing automatic network optimization as business needs and utilization change.

Access optimization products include:

- GeoSpan: for integrated, automated networked storage support for clusters
- Celerra HighRoad: for optimized delivery of files across SAN and NAS networked storage
- PowerPath: for dynamic multi-channel path optimization

Data Presentation abstracts network storage details and facilitates seamless information presentation and mobility. EMC will continue to develop and deliver products to further abstract or hide complex physical storage elements and present these storage components in simple business terms.

Data Mobility facilitates the simultaneous use and movement of information across networked storage without application interruption.

EMC is developing products that will enable the seamless movement of information throughout the storage infrastructure. These advanced tools will extend the value of information assets by enabling seamless re-location, re-purposing, and re-use of information in response to new business opportunities and changing imperatives.

Summary

Symmetrix systems are the most capable in the world. Their unequalled functionality rises to an even higher level with the introduction of the Symmetrix Direct Matrix Architecture.

The Symmetrix Direct Matrix Architecture runs Enginuity, a sophisticated, time-tested storage operating environment with functionality far beyond the simple microcode-based storage control of competitive products.

Enginuity provides stability and continuity across all Symmetrix generations. It fully exploits the performance potential of the Symmetrix hardware; ensures data availability; and sets a new standard of openness for broad integration and assured interoperability.

Networked storage from EMC unifies technologies, platforms, software and services across the enterprise, so information can be cost-effectively managed, protected, and shared. The Symmetrix DMX Architecture allows the storage for thousands of servers to be consolidated into a single storage array with dozens of terabytes of capacity, with uncompromising performance, availability, and functionality.

A Symmetrix environment can integrate SAN and NAS technologies to leverage the unique values of each.

Comprehensive interoperability testing ensures that technologies, software, and hardware components work together seamlessly, for flawless operation in the most complex environments.

EMC's *Business Continuity Software* provides industry-leading remote mirroring capabilities designed to minimize or eliminate the potential business impact of information loss or inaccessibility.

EMC's *Open Storage Management* software strategy extends Symmetrix functionality throughout a diverse, multi-vendor infrastructure. It affords system administrators unprecedented ease of management with a high degree of control.

Chapter Six: The Economics of Symmetrix DMX

A paradigm shift in the economics of high end storage

The Symmetrix DMX Architecture fundamentally changes the economics of high-end storage, opening up a formerly exclusive “club” to any IT organization faced with skyrocketing service levels.

Traditionally, entry into the high end has been an all or nothing proposition. Customers have been purchasing large systems, with capacities that often exceeded their immediate needs. Choice of systems has been severely limited, and the acquisition strategy has offered few options.

In addition, the systems themselves have required a climate-controlled, raised-floor environment, expensive to maintain and taking up a considerable amount of space.

In contrast, the incremental scalability and flexibility of the Symmetrix DMX Architecture lends itself to a range of packaging options.

The architecture opens up numerous possibilities in terms of acquisition strategy, so choices can be made based on the unique requirements of the environment.

The architecture also makes possible an unprecedented level of consolidation, for exceptional ease of management and the lowest possible TCO, with enormous capacity in a very small space.

And, the architecture allows Parity RAID and RAID 5 to run at very high performance levels, delivering a low cost-per-megabyte, with absolutely no performance degradation.

Incremental scalability: a key attribute of the Symmetrix DMX Architecture

The incremental scalability of the Symmetrix DMX Architecture allows configuration of the optimal amount of resources (Channel Directors, Disk Directors, Global Memory Directors, disk devices) to balance specific performance and scalability requirements with the cost of the system.

Incremental scalability is the key attribute that enables Symmetrix DMX series systems to scale up and down, so IT organizations can start with a basic system and add capacity, connectivity, and performance on a “buy what you need, when you need it” basis. For example, entry to high-end storage can be achieved with the rack-mounted Symmetrix DMX800 system and two disk array enclosures. Additional Disk Array Enclosures can be added as needed to a standard 19-inch rack.

Symmetrix DMX series systems

Symmetrix DMX series systems are the world's first storage systems based on a Direct Matrix Architecture. They are designed to provide unmatched levels of performance, connectivity, and scalability.

There are four models: the incrementally scalable DMX800, the single-bay DMX1000, the dual-bay DMX2000, and the triple-bay DMX3000.

Each system is offered in several targeted configurations designed to meet customer performance, connectivity, and economic requirements.

All Symmetrix DMX systems support the Enginuity storage operating environment, providing operational consistency across the entire Symmetrix family.


Symmetrix DMX800

The Symmetrix DMX800 is the industry's first high-end storage system contained in a standard 19-inch rack featuring a scalable design that allows Disk Array Enclosures to be added to scale capacity. It is especially attractive to cost-conscious IT organizations because space, cooling, and power requirements are modest in comparison with larger systems, making it economical to run and maintain.


The DMX800 can be configured with 8 to 120 drives for a total raw capacity of up to 17.5TB. It includes 8 or 16 2GB Fibre Channel drive ports; up to 16 2GB Fibre Channel I/O ports; or 8 2GB Fibre Channel ports with up to 4 FICON and/or Gigabit Ethernet or iSCSI ports. Global memory can be from 4 to 64GB. System expansion is economically achieved within the storage frame and by adding more disk expansion chassis, up to the limits described.

The EMC DMX multi-protocol channel director/adaptor, compatible with the DMX800 entry-level configuration and all other DMX models, provides unique connectivity configurability. It allows each of the two or four ports to be individually configured as FICON for mainframe connectivity, Gigabit Ethernet for remote mirroring over IP, or iSCSI.

The Symmetrix DMX800 is the solution of choice for customers who need the performance, availability, and functionality of a high-end storage solution, but have special deployment requirements or require less storage capacity.



Symmetrix DMX800




Symmetrix DMX1000

Symmetrix DMX1000

The Symmetrix DMX 1000 is a single-bay, high-performance, networked storage system. It has from 64 to 144 drives, for a maximum capacity of up to 21TB; 4 to 128GB of global memory; 16 to 48 2GB Fibre Channel host ports or ESCON channels; and/or up to 24 FICON, Gigabit Ethernet, or iSCSI connections.


The DMX1000 system is the ideal solution when a fully integrated storage system is required, such as for consolidation in heavy transaction processing environments.



Symmetrix DMX2000

Symmetrix DMX2000

The Symmetrix DMX2000 is a dual-bay configuration that accommodates more memory, more connectivity, and up to double the storage of the DMX1000. It features 128 to 288 drives for a maximum capacity of 42TB; 8 to 256GB of global memory; 16 to 96 2GB Fibre Channel host ports or ESCON channels; and/or up to 48 FICON, Gigabit Ethernet, or iSCSI connections.



Symmetrix DMX3000

Symmetrix DMX3000

Designed for the most demanding high-performance/high capacity consolidation requirements, the triple-bay DMX3000 features 289 to 576 drives for a maximum capacity of over 84TB, and usable capacity of up to 73.5TB. It offers up to 256GB of global memory and up to 64 Fibre Channel host ports or FICON channels and/or 32 FICON, Gigabit Ethernet, or iSCSI connections.

Symmetrix DMX optional configurations

A key characteristic of Symmetrix DMX series systems is the degree to which they can be custom configured to the precise needs of the environment. The following is a summary of options:

	Symmetrix DMX800	Symmetrix DMX1000	Symmetrix DMX2000	Symmetrix DMX3000
Packaging	Rack Mount (19")	Single Bay	Dual Bay	Triple Bay
Max Capacity (Raw)	8.76 TB (73GB disks) 17.5 TB (146GB disks)	10.5 TB (73GB disks) 21.0 TB (146GB disks)	21.0 TB (73GB disks) 42.0 TB (146GB disks)	42.1 TB (73GB disks) 84.0 TB (146GB disks)
Connectivity	8-16 2Gbit ports	16-48 2Gbit ports	16-96 2Gbit ports	16-64 2Gbit ports
ESCON Directors	N/A	16-48 ESCON ports	16-96 ESCON ports	16-64 ESCON ports
FICON Directors	2-4 FICON ports	4-24 FICON ports	8-48 FICON ports	8-32 FICON ports
Max Disks	120	144	288	576
Disk Directors	2	2-4	4-8	8
Disk Channels (Fibre Loops)	8 Redundant Channels	8-16 Redundant Channels	16-32 Redundant Channels	32 Redundant Channels
Max Disk Drives per Disk Channel	15	9-18	9-18	18
Cache	4-64GB	4-128GB	8-256GB	8-256GB

The economics of acquisition

With Symmetrix DMX, the high-end storage acquisition strategy is chosen, not imposed.

A primary advantage of buying into the Symmetrix DMX Architecture is the ability to purchase precisely what is needed, when it is needed. This advantage is a function of 1] the modularity of the architecture itself, and 2] the fact that it plays in a networked storage environment.

As an example, if the immediate need is for 5TB of e-mail storage, it is now practical to buy only those 5TB instead of, as in the past, buying 15 TB of consolidated storage and then "growing into it" over time.

In that example, the immediate requirement could be filled by a rack-mounted, modular Symmetrix DMX800. If, subsequently, an additional 5TB are needed, that requirement could be economically met by acquiring another DMX800.

As capacity requirements grow, the answer might be additional DMX800 systems, or, perhaps, a single DMX2000. The point is that the modularity and networking capabilities of the architecture, plus the affordable cost of the DMX800, open up an entire range of acquisition options that never existed before.

The economics of consolidation

With the Symmetrix DMX Architecture, an unprecedented degree of consolidation is now possible. Consider that the footprint of single-bay Symmetrix DMX1000 system is a single floor tile. Whereas in the past, a single-bay cabinet could hold up to 96 disk drives, the DMX1000 accommodates up to 144—a significant capacity increase with no increase in required floor space.

In addition to raw capacity, Symmetrix DMX1000, DMX2000, and DMX3000 systems deliver twice the server connectivity—again, in the same space. The ability to connect more of the business into a single storage array saves money. Consolidated storage is easier to manage and delivers far higher service levels across the entire enterprise.

In addition to its affordable price, the sheer performance of the entry-level Symmetrix DMX800 brings unexpected economic benefits. This small, rack-mounted system actually outperforms massive machines from other vendors—machines that not only take up a lot of space, but have expensive environmental requirements as well.

The economics of RAID 5 and Parity RAID

RAID 5 and Parity RAID provide a cost-effective means of protecting data against drive failure. While the most demanding environments continue to opt for mirrored storage for maximum performance, RAID 5 and Parity RAID offer an extremely attractive alternative for information storage offering strong performance and economics. They offer far better use of disk resources, while minimizing the amount of hardware that has to be purchased, managed, and maintained.

Summary

The introduction of the Symmetrix DMX Architecture represents a paradigm shift in the economics of high-end storage.

The modularity and flexibility of the architecture open up entirely new possibilities in terms of acquisition strategy, packaging options, configurability, connectivity, scalability, and ease of consolidation.

The architecture also delivers unprecedented high-end performance at a surprisingly low cost-per-megabyte.

Symmetrix DMX systems range from the rack-mounted Symmetrix DMX800, with total raw capacity of 17.5TB, to the triple-bay Symmetrix DMX3000, with a maximum capacity of 84TB. Symmetrix DMX systems are offered in several configuration options designed for even the most demanding application environments.

By design, Symmetrix DMX systems can be readily custom-configured to the precise requirements of the environment.

The Symmetrix DMX Architecture's modularity and networking capabilities make it easy to purchase only the storage that's needed, when it is needed, and scale up incrementally as requirements grow.

The Symmetrix DMX Architecture facilitates consolidation. It packs enormous disk density and connectivity in the least possible floor space. And, even the smallest of the systems, the Symmetrix DMX800, outperforms competitive machines that are much larger and that impose expensive environmental requirements.


With both Parity RAID and RAID 5, Symmetrix DMX makes the most efficient possible use of disk resources, while minimizing the amount of hardware that has to be purchased, managed, and maintained.

Chapter Seven: Services

EMC Services

In addition to uncompromising performance, availability, functionality, and economic benefits, Symmetrix DMX users also have access to the world's top-rated customer service and support.

The EMC Global Services organization helps the enterprise make the most of its EMC technology investment with a continuum of best-in-class services spanning the entire information lifecycle.



Information Solutions Consulting: Evaluating Change

Information Solutions Consulting (ISC) specializes in helping IT organizations determine how to best deploy their information resources to meet changing business needs.

Through its comprehensive, strategic consulting approach and proprietary financial modeling tools, ISC delivers the following offerings:

- *Storage Infrastructure Strategy* aligns the information storage environment with long-term business goals. EMC assesses current information practices and makes recommendations for improvement.

- *Storage Management Optimization* controls costs and complexities associated with information growth and storage utilization. EMC identifies ways to make better use of current storage infrastructure while planning for change.
- *Information Storage Consolidation* makes the most of current and future investments by building a more efficient, manageable infrastructure. EMC helps the enterprise develop and implement a plan to consolidate storage technologies, increasing efficiency while controlling or reducing costs.
- *Business Continuity Planning* helps prepare for unexpected interruptions and planned downtime with best practices that keep the business—not just the data—running smoothly. This engagement covers the development and testing of the plan, as well as guidance on its ongoing refinement.

EMC Services: Planning, Implementing, and Initiating Operations

EMC consultants deliver the business benefits of Automated Networked Storage to the enterprise through the following services:

- Networked Storage brings DAS, NAS, and SAN environments together in one unified infrastructure for maximum information availability and ease of management.
- Open Storage Management simplifies the management of open storage with EMC's Automated Information Storage strategy and software.
- Business Continuity includes design and implementation of a business continuity plan to help ensure uninterrupted system and data availability. The plan leverages advanced replication and mirroring strategies made possible by the new Symmetrix DMX Architecture and state-of-the-art EMC software.

Proven Methodologies

Every EMC Global Services engagement is guided by a proven methodology, the Global Delivery Model (GDM). The GDM assures rapid, flawless implementation in every engagement around the world.

EMC is committed to complete customer satisfaction. To this end, all EMC Global Services technologists undergo comprehensive training and certification in the industry's most advanced storage technology and implementation methodologies.

As the leader in Automated Networked Storage, EMC has the expertise and project management skills to ensure maximum value and minimal disruption during any networked storage engagement.

Customer Service: Ongoing Operations

EMC Customer Service delivers global, proactive support with a network of over 5,000 technical, field, and support personnel. EMC's pre-emptive approach to support means problems are addressed and often eliminated before they occur. Remote support provides notification when there's a problem or a potential problem. If onsite service is required, field staff can hot-swap a part without system downtime. EMC parts depots are accessible 24x7.

Remote Support

EMC Symmetrix systems are equipped with automatic phone-home capabilities, so EMC service experts can monitor a system 24x7. And by dialing back into the EMC system, they can take action quickly, analyzing events and abnormalities, and resolving most issues before they affect business. Advanced remote support means a proactive and pre-emptive approach unmatched in the industry.

Software Support

An all-inclusive software support and maintenance program ensures optimum availability of mission-critical information. EMC software specialists provide 24x7 telephone support to meet the needs of the most complex multi-vendor environment. And EMC e-services like Powerlink and Knowledgebase make information, solutions, and software upgrades instantly accessible.

Change Control

EMC's industry-leading change control process leverages the outstanding connectivity, flexibility, and upgradeability engineered into every Symmetrix DMX system. EMC experts meticulously plan and orchestrate changes to the EMC solution, from standard microcode upgrades to massive data center relocations.

Installation Support

EMC specialists configure and install Symmetrix DMX systems according to customer specifications and business requirements. They create file systems and set access rights, as required; export file systems to the network; mount file systems on individual machines; and provide channel and network connectivity.

Post-sale Warranty and Product Support

Post-sale Warranty Coverage of Symmetrix DMX systems includes EMC's basic two-year hardware and 90-day software warranty plan, with 24x7 coverage. Post-warranty service offerings include 24x7 coverage, technical support, and service and maintenance contracts.

Worldwide Organization, Local Support

The EMC Customer Support Center, headquartered in the United States, directly supports EMC hardware and software products. Use the following numbers to contact EMC and obtain technical support:

U.S.: (800) 782-4362 (SVC-4EMC)

Canada: (800) 543-4782 (543-4SVC)

Worldwide: 1 + (508) 497-7901 (or contact the nearest EMC office)

Global Technical Training

EMC Global Technical Training delivers ongoing technical education that gives customers the knowledge they need to maximize the value of their technology investment for the best competitive advantage. Both e-learning and traditional instruction are available. EMC E-learning incorporates online learning into the suite of training, education, and certification solutions available to customers, partners, and employees.

The EMC Proven Professional Certification Program

The EMC Proven Professional Certification Program is aligned with other IT industry certification programs, notably Microsoft and Cisco. Students can achieve an Associates or a Masters level of certification. Four tracks are offered, based on IT job roles: Operator (data center operations management), Builder (datacenter implementation and integration), Architect (enterprise networking solution design), and Instructor (networked storage infrastructure knowledge transfer).

EMC Powerlink

EMC Powerlink showcases EMC products and services at work in an interactive portal, providing invaluable insights and a wealth of practical information. EMC Powerlink delivers continuous value for customers and partners with:

- a 24x7 connection to product and technical information
- online services and support
- training and certification programs
- collaboration with product specialists

Chapter Eight: Summary

Summary

The Symmetrix DMX Architecture is a high-end storage technology designed for the exponentially rising service level demands of the next decade. It delivers uncompromising performance, availability, and functionality, while fundamentally changing the economics of high-end storage. To recap some salient points:

- The Symmetrix DMX Architecture is a significant advance over switch and bus-based architectures. Direct access from the front of the storage array to the back totally eliminates contention issues, provides unlimited scalability, and creates a system with extraordinary modular flexibility.
- Sustained, burst, and business continuance performance of the Symmetrix DMX Architecture is unmatched in the industry. Performance in a Symmetrix DMX environment is a function of the accumulated algorithmic intelligence of the Enginuity storage operating environment and the direct matrix architecture.
- Symmetrix DMX delivers unequalled availability, with the world's most fault-tolerant design, featuring full redundancy, proactive monitoring, and error detection and correction. Symmetrix DMX systems offer several disk protection options including mirrored disk, RAID 5, Parity RAID, and Remote Data Facility products.
- The proven functionality of Symmetrix systems attains its highest level ever with the introduction of the Symmetrix DMX Architecture. And, the time-tested Enginuity storage operating environment provides absolute functional continuity with previous Symmetrix generations.
- Symmetrix DMX heralds a paradigm shift in the economics of high-end storage, putting high-end functionality within reach of virtually any enterprise facing skyrocketing service level requirements. Packaging options range from high-performance entry-level systems, to specially configured machines for the most extreme-demand storage environments.
- In addition to uncompromising performance, availability, functionality, and economic benefits, Symmetrix DMX users have access to a customer service organization ranked #1 in any industry. Support is proactive, comprehensive, 24x7, and global.

Symmetrix DMX and Automated Networked Storage

The introduction of the Symmetrix DMX Architecture has fundamentally altered customer expectations in terms of availability, performance, replication, scalability, and management for the high-end storage environment.

To users of Symmetrix DMX series systems, high-end availability is more than just redundancy—it means non-disruptive operations and upgrades, and “always

online.” High-end performance means handling all workloads, predictable or not, under all conditions. High-end replication means copying any amount of data any-time and sending it any distance. High-end scalability means more than capacity—it means having the flexibility to handle any service level or application, cost-effectively, no matter how the business changes. And, high-end storage management goes beyond monitoring storage arrays—it means managing the service levels the business expects, from provisioning to business continuity.

Symmetrix DMX has redefined the standard in high-end storage. However, Symmetrix DMX represents just one component of EMC’s Automated Networked Storage (ANS) strategy.

Automated Networked Storage is an integrated networked storage infrastructure that combines storage area networks, network attached storage, and content addressed storage, all managed by the most advanced open storage management software in the industry.

ANS is predicated on the belief that the whole is more than the sum of its parts—that ultimate value is achieved by unifying storage networking technologies, storage platforms, and services.

In a complex, multi-vendor world, the payoff to the enterprise is the power and control to cost-effectively manage, protect, and share information by:

- Consolidating physical infrastructure
- Centralizing or consolidating management and people
- Achieving absolute business continuity, no matter what
- Having the assured flexibility to support changing service level requirements

EMC believes that Symmetrix DMX is the most powerful high-end storage architecture ever devised for leveraging the full potential of Automated Networked Storage.



EMC²

where information lives

EMC², EMC, Symmetrix, Celerra, CLARiiON, CLARAlert, DG, E-Infostructure, HighRoad, Navisphere, PowerPath, ResourcePak, SnapView/IP, SRDF, TimeFinder, VisualSAN, where information lives, and The EMC Effect are registered trademarks and EMC Automated Networked Storage, EMC ControlCenter, EMC Developers Program, EMC Enterprise Storage, EMC Enterprise Storage Network, EMC OnCourse, EMC Proven, EMC Snap, Access Logix, AutoAdvice, Automated Resource Manager, AutoSwap, AVALONidm, C-Clip, CacheStorm, Celerra Replicator, Centera, CentraStar, CLARevent, Connectrix, CopyCross, CopyPoint, CrosStor, Direct Matrix, Direct Matrix Architecture, EDM, E-Lab, Enginuity, FarPoint, FLARE, GeoSpan, InfoMover, MirrorView, NetWin, OnAlert, OpenScale, Powerlink, PowerVolume, RepliCare, SafeLine, SAN Architect, SAN Copy, SAN Manager, SDMS, SnapSure, SnapView, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix DMX, Universal Data Tone, and VisualSRM are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners.

©2003, 2004 EMC Corporation. All rights reserved. Printed in the USA. 03/04
Product Description Guide
C1011.1