



Strategic Snapshot

EMC Drives the High End to New Frontiers

By Joyce Tompsett Becknell

The Sageza Group, Inc.
July 2005

sageza.com
info@sageza.com

The Sageza Group, Inc.
32108 Alvarado Blvd #354
Union City, CA 94587
650-390-0700 fax 650-649-2302
London +44 (0) 20-7900-2819
Milan +39 02-9544-1646

EMC Drives the High End to New Frontiers

ABSTRACT

Storage is one of the fastest growing segments of the IT infrastructure. The data it houses and makes available is critical to business success. New products in the market provide superior capabilities that allow companies to better manage and control their information assets; however, it is difficult for most companies to take strategic advantage of those features. Often companies don't fully understand their data and have not implemented an architecture that aligns it, classifies it, protects it, and enables the business to leverage it. EMC is now rising to these challenges on a couple of fronts. Along with a new Symmetrix DMX-3 for high-end storage needs which advances the capabilities of high-end storage, the company also has new EMC Consulting Services. These offerings provide customers with the tools to better understand their data and storage environment, such as classification and policy, or architecture and consolidation services. With the DMX-3 in combination with a strategic understanding of their corporate data, customers are better-positioned to take advantage of the resources they already have, add new ones for ongoing storage development, and strategically drive their businesses forward with superior information lifecycle capabilities.

EMC Drives the High End to New Frontiers

TABLE OF CONTENTS

Architecting the Future of Enterprise Data 1

EMC Makes the High End ILM Capable 1

Services to Help Enterprise Get Started..... 3

Conclusion: Bringing Strategic Value to ILM 4

Architecting the Future of Enterprise Data

Information is made up of underlying business data in context, such as customer relations, the retail process, corporate governance, or financial statistics. In order to manage information, many vendors have positioned their products within the context of information lifecycle management (ILM). This is the principle of the first phase of ILM: putting data on the right type of storage over its lifecycle, so that it sits in the most cost-effective location at any point in time, with the right amount of protection, and retrievable in an appropriate window. Product features and capabilities can be best understood in light of ILM, and many new products are coming to market within a specific ILM-driven context. This phase of ILM is sensible and easy to grasp for most IT and business managers and well on its way to near-universal adoption. Many organizations are moving to the next phases of ILM; that is, aligning and automating this infrastructure for a specific application (i.e., email or content management) and then across the infrastructure and across applications. To move towards these later phases of ILM, a few organizations are beginning to develop corporate data architecture and an understanding of how data flows among business processes and applications. This understanding is important as companies leverage their existing infrastructure, add new elements, and use their storage budgets most effectively to grow the top line, improve the bottom line, and improve IT efficiencies.

EMC has been a proponent of ILM solutions for a while now, positioning its products and services to fit the needs of an evolving information lifecycle. EMC understands that organizations have a range of information needs, from environments where high performance is critical to environments that scale rapidly, such as retail transactions, and even information that needs to be preserved for extended periods of time. With this understanding in mind, EMC is announcing a new high-end storage product, the DMX-3, which continues to offer features designed to make complex information environments more manageable with cost-effective technology. For customers who are still sorting out their data architecture, EMC has also developed services to help customers evaluate their existing infrastructure and data needs. In this paper we look at the new DMX-3 and how EMC is driving their high-end storage offering into new territory, and look at some of the services customers can use to take better advantage of all that the new system has to offer.

EMC Makes the High End ILM Capable

The information explosion continues, and companies are taking advantage of the latest technology advancements to improve their internal infrastructure and to provide new products and services to customers. At the same time, the proliferation of data combined with new technologies has a tendency to increase underlying architectural complexity. IT managers are realizing the strategic value of data for corporate information and wanting to implement appropriate ILM strategies for their organizations. IT managers are shifting their focus from getting their data into the underlying storage and protecting and ensuring its availability, to getting the data/information out of the storage, in business and application context, to drive business growth, differentiation, and innovation. EMC has designed the latest version of the DMX architecture, the DMX-3, to address these changing high-end requirements. EMC is augmenting its leading-edge technology with more capabilities and features designed to address the qualities it believes are most important for its customers.

EMC believes that the most important capabilities needed in high-end systems now are the usual increases in scalability and performance. EMC continues to drive these requirements into its products, combined with capabilities that make ILM easier to implement with less complexity and improved reliability, while preserving the investment customers have already

made by leveraging software compatibility with previous versions of the DMX family. For the increased scalability and performance, EMC has taken several steps. Most directly, the product has been made bigger. If the current DMX models scaled from 30 to 576 drives, the new DMX-3 scales from 360 to 960 drives, nearly doubling the capacity of the former (with announced plans to support 1,920 drives by mid-2006 and more than a Petabyte once new low-cost Fibre Channel drives are available and qualified in 2006). Beyond size, EMC also believes that scalability implies flexibility and the ability to provide predictable performance in an unpredictable world. That means the ability to change technologies as needed, the ability to change without disruption to users, and the ability to change the underlying storage as business requirements change. EMC also believes change should be modular, and that customers should be able to scale either capacity or performance, allowing them to adapt their storage capabilities to current business needs.

One of the fundamental parts of an ILM strategy is tiered storage. This is the idea that more important data should sit on the more expensive higher-performing storage, and that as data becomes less important it should be moved to less expensive, slower storage to help lower costs and free up resources for more strategic use. This has usually meant that data is physically moved from one storage array to another storage array, requiring the purchase of additional storage or additional capacity and physically migrating storage from one system to another, all of which are opportunities for encroaching complexity. With the DMX-3 customers can now do tiered storage within one box for various levels of fibre channel storage. For low capacity but high performance, EMC has 73GB and 146GB drives that run at 15k RPM. For medium capacity and medium performance, EMC has 146GB and 300GB drives that run at 10k RPM. For high capacity but lower performance, EMC will support the emerging/forementioned low-cost Fibre Channel 400GB and 500GB drives that run at 7200 RPM. These drives bring \$ per GB levels approaching that of ATA technology to the high end, enabling a high-end array to support so-called "tier 2" applications economically. EMC also notes that, significantly, the low-cost FC drives will also be supported on the entire DMX line, not just the DMX-3.

EMC has several ways to make management easier, which the company believes is another prerequisite of high-end storage in an ILM world. Manageability software used to mean the ability to set up the array and occasionally add capacity to it. Now manageability means the ability to shift resources dynamically as the business requires, without disruption to users and applications. EMC has addressed this issue on several fronts. Symmetrix storage is designed to handle data from open systems as well as for mainframes. For open systems data, EMC has the Open Migrator/LM that can dynamically migrate volumes online without interrupting the database and allows block or file-based mobility. Open Migrator/LM provides SAN host-based relocation, so users should see no change in response times, making open systems data management more effective.

For mainframe systems, EMC has the Logical Data Migration Facility (LDMF). The LDMF (developed and marketed jointly with Softek) allows online migration and aggregation of mainframe volumes at the dataset level. It can combine smaller volumes into larger ones and update catalog entries online to reflect new dataset locations. With LDMF customers can consolidate volumes onto larger arrays, reducing the number of logical volumes to be managed. LDMF also helps to reduce the UCB address space requirements, as MVS has a limit of 64k UCBs which can otherwise limit consolidation and tiered storage scalability. In practice, EMC believes the DMX is the first tiered storage solution for the mainframe space, allowing mainframe customers to consolidate their smaller volumes into larger ones. In conjunction with the protection and availability benefits of Symmetrix for all tiers, including replication consistency across interdependent tier 1 and tier 2 data, managers can now achieve significant synergies in business continuity/disaster recovery plans. The ability to

mix these drives within the array makes it fundamentally easier to tier and manage high-end storage.

Engenuity is the operating environment for the Symmetrix family, and the foundational base for other EMC software. With the announcement of the DMX-3, EMC is also adding numerous enhancements to Engenuity and some of the software that sits on top of it including SRDF, the remote data facility; and TimeFinder, EMC's local storage replication product. For the DMX-3 in particular, EMC will add support for a faster Direct Matrix (the DMX array architecture) and mirrored DDR-based global memory (expandable up to 512GB) with power fail vault protection. Each channel board now has eight processors on it (two per "slice") with each slice controlling two ports. In addition, Engenuity will support in excess of 2,000 drives per array, 255 hosts per fibre channel port, 512 hosts per SCSI port, up to 32,000 logical devices (64,000 supported in the near future), and for the mainframe customer, up to 31 Parallel Access Volume (PAV) aliases (255 PAVs supported in the near future).

For customers who already have the second-generation DMX, there may be some question as to whether they should move to DMX-3 or where to add it within their storage infrastructure. The DMX-3 is designed for very large-scale consolidation. It is larger than a DMX3000, and it is capable of consolidating multiple applications, tiers, or arrays into a single large array. If customers have a need for ever-larger arrays, then the DMX-3 is good for reducing the TCO by reducing the hardware and software acquisition price and providing lower costs for management over time. The lower costs come in a couple of ways. First, the DMX-3 is faster than a DMX-2 with the same number of spindles and DAs. In essence EMC estimates it provides twice the IOPs for only 20% to 30% greater price. Second, it is more cost-effective for customers who are thinking about deploying multiple DMX3000s. EMC estimates that one DMX-3 system with 960 drives is 10% to 20% less than two DMX3000s with about 500 drives each. The heart of EMC's design goal for the DMX-3 was performance at capacity. This system was designed from the ground up to handle some of the largest data processing environments, while delivering predictable service levels and enterprise-class functionality.

Services to Help Enterprise Get Started

While some customers may know exactly how they want to deploy the DMX-3 technology within their organization, others may want to use EMC to help them do more than manage the physical location of their data. They may want help tying architectural changes to business needs, or they may need a more basic understanding of where data sits in their architecture and how to better manage it. To that end, EMC has announced new EMC Consulting Services, which are designed to optimize infrastructure and convert IT to a strategic business asset by aligning technology to business needs. EMC has classified the services into four groups: Classification and Policy Services, Architecture and Consolidation Services, Storage Management Optimization Services, and Information Protection Services. To take advantage of the capabilities of the DMX-3, some IT managers may find that Classification and Policy Services or Architecture and Consolidation Services may be their best bet. Classification and Policy Services can help customers classify their information and data and help them with storage tiering, with policy-based information management, and with compliance solutions. Architecture and Consolidation Services take this one step further by helping customers design a consolidation strategy driven by business requirements or with more generic storage and server consolidation. With these services EMC helps the customer take maximum advantage of the technology offered within the context of their current and future infrastructure and helps them reach their strategic business goals.

Conclusion: Bringing Strategic Value to ILM

Information lifecycle management requires environments that can handle complex data needs that change with time. The proper care and feeding of data requires that systems be scalable, modular to grow where growth is needed, cost-effective, and easier to manage, lowering administrative costs and freeing administrators to perform more strategic tasks rather than firefighting. EMC's new DMX-3 is the next step in the evolution of high-end storage that seeks to be less monolithic and more adaptable to complex, dynamic environments. A capability such as tiering within one system rather than across arrays is one example of how the DMX-3 can make ILM easier for companies to achieve. The scalability and flexibility directly translate into opportunities for consolidation of many smaller arrays into one larger array, reducing management complexity. Finally, improved data migration and mobility services can support mission-critical data in the appropriate manner for both open systems and mainframe data.

At the same time, for customers who are still working on turning storage into a part of their strategic data architecture, EMC now offers services that allow customers to classify and therefore better understand what they have in terms of data and physical resources and map where they are today against where they ultimately want to go. These services can help customers take better advantage of the new features of the DMX-3 and help them understand how to deploy these systems for greatest strategic improvement. Customers who are aligning their IT infrastructure to their business will find the EMC products and services offer many of the tools to make the data and storage infrastructure align with ILM and demand driven business needs.