

STK Announces New ILM Initiative — Automating the Storage and Management of Data

Analyst: David Reine

Management Summary

The theme is the same in every data center. From the small and medium business at the low-end of the storage spectrum, to the largest enterprises at its high-end, **the effort to simplify storage access and control costs is on the rise.** Despite the fact that the cost per gigabyte of disk capacity continues to fall, the volume of information and the costs associated with its management continue to rise. For example, the industry is looking at an increase in the number of business emails per employee from about 75/day in 2003 to a projected 100 in 2007. This represents a projected increase in storage of from about 7 to 15 MB per day. With new regulations requiring the long-term preservation of email, an enterprise of 1000 employees will increase storage by over 300GB per month. E-mail servers cannot store such massive amounts of data for long periods. While information technology (IT) managers recognizes the short-term value of storing email, they simply have not been able to cope with the long-term value of that email changing over time. **The concept of Information Lifecycle Management (ILM) enables the enterprise of today to meet the challenges of today that require balancing the cost of storing and managing information with its changing value over time.**

Another example in the medical industry is the hospital data center. There is an urgency to keep a patient's latest MRI in digital format on the enterprise-level storage array for instant access while the patient is in the hospital. However, once the patient leaves the hospital, the data center may move that image to a secondary storage device when immediate access is no longer required. This image is not required again until the patient's check-up, perhaps a year later. Meanwhile, gigabytes of primary storage can be freed to capture the latest images for current patients, eliminating the need to add additional drives. Policies need to be established that automate this function. In order to be useful, however, ILM must be more than just a concept. It must have a practical methodology for aligning storage costs with business priorities. ILM needs to be able to enhance an enterprise's ability to manage the information as it changes over its lifetime. **ILM needs to be able to automate the movement and management of data, providing alternatives for storing the data at each lifecycle stage.**

StorageTek (STK) has now introduced a series of tools to simplify information management and to automate the movement and management of that information throughout its lifetime. In order to learn more about STK's approach to ILM, please read on.

IN THIS ISSUE

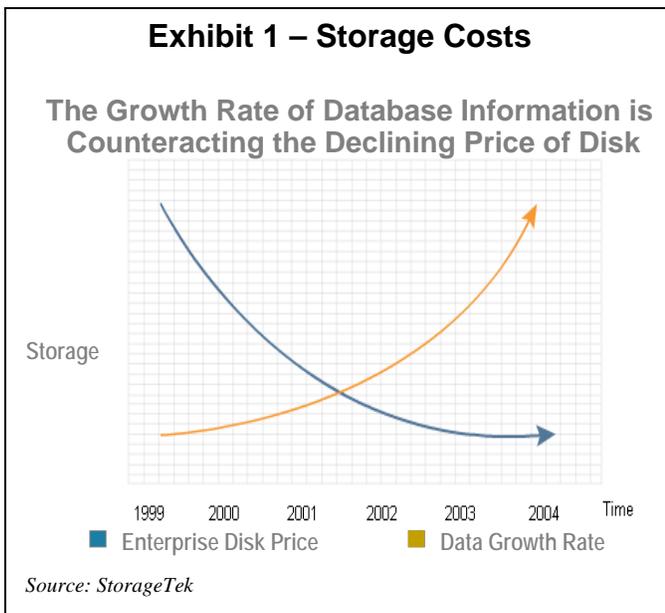
- **The Data Center Storage Morass** 2
- **The StorageTek ILM Solution** 2
- **Conclusion** 5

The Data Center Storage Morass

As we proceed into 2004, enterprises continue to gather more information that is essential to attaining the corporate goal of profitability. As business activity increases, disk capacity rises; the requirements for additional storage stress the budgets of information technology operations even further than they already are. Despite the fact that the cost of disk devices continues to fall (see Exhibit 1, below), the increase in the number of devices required wipes out any potential savings. Furthermore, the costs associated with the protection and management of that data rise even more. Data centers have to find new and inventive ways to improve the storage mechanics of their data, keep it safe, and efficiently retain it over time. Moreover, they must do this within the limits of a budget, while learning to adhere to an entirely new set of industry and government standards, instituted for the preservation of financial and historical data.

growth in all of these areas is noteworthy, it is even more important to realize that the vast majority of this information is no longer useful 30 to 60 days after creation. Take order processing, for example. A customer places an order. Manufacturing builds it, Traffic ships it, Accounts Payable bills it, and Accounts Receivable collects it. This entire process should take less than 60 days. However, because of accounting standards, all of these transactions will remain on the enterprise's primary storage facility, consuming the most expensive form of storage in the data center, unless IT implements some form of information lifecycle management.

Today, the data center is looking to reduce the overall cost of storage while at the same time implementing an efficient process to access required data throughout the life of that data. This includes improving the data protection policy for backup and recovery on a daily and weekly basis, as well as the disaster recovery policy for emergency reconstruction. Further, they are looking to accomplish this within the guidelines of regulatory compliance. **StorageTek is responding to these needs with more than just a concept for ILM. They have implemented a full-reference architecture with turnkey solutions for the management of the total cost of storage ownership.** These solutions include focused applications and multi-tiered hardware to manage the cost of the storage against value of the data.



One of the areas of greatest storage growth is in databases. Typically, databases consume 50-75% of enterprise storage capacity. Significantly, industry reports state that these databases are growing at an annual rate of 125%. Obviously, this does not include the growth in email and other digital-based structures, such as movies, medical images, and order catalogs. While the

The StorageTek ILM Solution

In STK's view, **ILM involves balancing the cost of storing and managing information with its changing value over time.** Their solution can be broken down into three categories:

1. Reference Architectures,
2. Storage Software Solutions, and
3. Storage Hardware products.

Let us look at each of these categories.

Reference Architectures

StorageTek has implemented a series of

reference architectures covering a variety of open systems and mainframe requirements. All of these architectures demonstrate how STK applies ILM to solve critical information management issues. They include multiple **STK and third-party products**, integrated and validated in StorageTek engineering facilities. The data center can integrate these architectures into existing storage infrastructures, although they may substitute other manufacturer's products:

- Open Systems Archiving
- Mainframe Archiving
- Open Systems Data Protection

When your responsibility for data is outgrowing your ability to manage it, these architectures provide examples of how ILM can be a reliable solution, using best practices, to gain control of all of your data. StorageTek has also put in place the following set of focused services to ensure the reduction of the total cost of ownership (TCO) related to storage architecture:

- Storage Appraisal
- Backup/Recovery Optimization
- Disk Capacity Optimization
- SAN Optimization
- Storage Management Gap Analysis

Open Systems Archiving

This architecture delivers a complete set of cost-effective archive solutions that automatically manage data across a tiered storage environment of open systems servers. These applications enable the establishment of data retention periods and the archiving and recall of mission-critical data from secondary fibre channel networked peripherals. Archiving applications are available for e-mail, medical imaging, video surveillance imaging, and regulatory compliance. These solutions target out-of-control archiving budgets and inadequate policies for storage management and data retention.

Mainframe Archiving

The mainframe architecture assists enterprises manage information cost-

effectively, across a tiered storage hierarchy in a mainframe environment while delivering identical solutions as in the Open Systems archiving arena. While the Open Systems archive uses fibre channel to *L180/L700* libraries, the mainframe archive architecture uses ESCON or FICON to connect to PowderHorn silos.

Open Systems Data Protection

This open systems architecture provides the enterprise with the protection assurance that its mission-critical data needs at all times, using fewer resources. It includes applications for backup and recovery, data replication, and disaster recovery in a Fibre Channel network to reduce the implementation time and risk often encountered in implementing complex data protection policies to protect the right data in the right way.

It enables policies to implement the new regulations that mandate the retention of more data for longer periods. It also enables the ability to establish policies to delete inappropriately retained privacy data while at the same time eliminating the problems encountered when older data becomes trapped in obsolete platforms.

ILM Software Solutions

ILM is STK's strategy for creating an adaptive storage hierarchy that dynamically aligns the value of the information with the appropriate tier of storage to improve the enterprise operation, while maintaining regulatory compliance. In this manner, STK can dynamically allocate storage resources to meet information requirements. At the same time, they can reduce the TCO of storage without affecting performance. The initial offerings in this area are:

- *STK Email Xcelerator* – a suite of applications that provide enterprises with an end-to-end e-mail archiving solution; and
- *STK Lifecycle Director for DB2* – software that dynamically manages expensive disk by moving low-access data to low-cost media.

StorageTek Email Xcelerator

Made available through a partnership with IXOS, *Email Xcelerator* maximizes the value of messaging data assets by offering the right solution to match a variety of business needs. It consists of three distinct solutions: *ArchiveMaster*, *ReportMaster*, and *AuditMaster*.

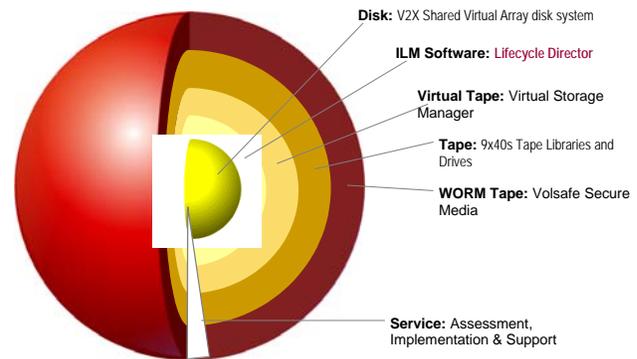
- *ArchiveMaster* – enables the enterprise to improve the e-mail lifecycle by migrating older messages and attachments to inexpensive media, saving primary storage space and reducing online storage costs.
- *ReportMaster* – allows the user to access the knowledge contained in the e-mail archive through advanced search and retrieval techniques.
- *AuditMaster* – enables regulated organizations to comply with policies pertaining to review, retention, and classification of messages and attachments.

With IT administration spending 8-12 hours/week on email backup and archiving, and another 5-6 hours recovering archived messages, Xcelerator creates a complete end-to-end solution when combined with STK's disk, tape, networking, Application Storage Manager, and services.

StorageTek Lifecycle Director for DB2

Lifecycle Director extends ILM principles to mainframe enterprises running DB2 databases. It dynamically cleans up expensive disk devices by moving aged-data to low-cost media, automatically and transparently, providing an end-to-end solution including hardware and software (See Exhibit 2, on the right). Therefore, the data center can free up extensive amounts of disk space, shrink the length of the backup window, and still maintain virtually instant access to the aged data. Thereby, they can extend the life of existing disk drives, reducing the requirement for additional disk acquisitions.

Exhibit 2 – Database ILM Solution



Source: StorageTek

The reduction of on-line storage also improves the performance and simplifies the maintenance of the system, with backups completing significantly faster while retaining almost instant access to archived data.

Tivoli Storage Manager

A third party management solution from IBM, *Tivoli Storage Manager* provides expanded options for storage management within the Open Systems data-protection reference architecture. It complements the established relationships that STK has had with *Veritas* and *Legato* in the area of backup and recovery management.

Storage Hardware Solutions

As we can see in Exhibit 2 below, STK offers a wide range of storage products to supplement its ILM strategy. They range from disk systems to disk arrays, from a simple tape drive to an extensive silo with the capability to support WORM tape. For purposes of this review, we will examine the newest entry into STK's disk storage offering, the *BladeStore B220* and *B280* disk systems and updates for the *T9840C* tape drive.

BladeStore B-Series

Complementing the *V-Series* virtual disk arrays and the *D-Series* disk systems from StorageTek, the *BladeStore B220* and *B280* disk systems present a new level of price/performance in the multi-tiered ILM

environment. These systems are based upon the widely accepted 2Gb Fibre Channel architecture of the *B220* and *B280 Controllers* and the *B250* disk array. **The B220 disk system can scale from 6.25TB of low-cost ATA disk to 75TB, while the B280 can scale to 150TB.** The *B250* disk array joins the *B200* SATA array in the STK arsenal of solutions.

In addition to their ability to offer low-cost, high-capacity storage for capacity-centric environments, the *B220* can also perform more data replication activities inexpensively, eliminating the need for expensive redundant hardware and software at a remote site. The *B220* can also be inserted between primary disk storage and tape to relieve the pressure on the data center concerning tape backup windows in a D-D-T environment. It will also improve the reliability of the backup while accelerating the recovery time. The *B220* can also replace costly optical disk and obsolete tape solutions to retrieve e-mail, medical images, and other fixed content storage.

The scalability and enterprise software of the *B280* makes it ideal for workgroup consolidation and use in online storage environments. These require mirroring and path failover but may not require the performance of a more expensive enterprise solution. This scalability enables the data center to implement a very cost-effective storage strategy with a “pay as you grow” implementation. This can fulfill any on-demand requirement established by IT management in the area of just-in-time disk availability to take advantage of falling disk prices.

T9840C Tape Drive

In order to satisfy the requirements of an ILM environment, any secondary storage device must be able to provide fast data access in a capacity-centric environment. With FICON and ESCON interfaces to improve the performance of the T9840 tape drive, STK can deliver a high-performance solution in zOS applications in which tape is an extension of disk and as a highly reliable data archive medium.

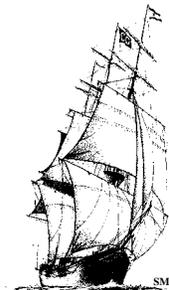
Mainframe archiving provides the environment for the T9840C to excel. With an average data access of 12 seconds, the T9840C fills the unique performance niche between disks with millisecond response time and capacity-centric tape drives with access times in the 40-85 second range. The faster transfer rates also mean more bandwidth to shrink the tape processing window with high duty-cycle reliability to ensure availability and a safe storage environment.

Conclusion

STK has simplified information management with new ways to automate the movement and management of data throughout its lifecycle. They have demonstrated value with ILM reference architectures that enable the automation of information management. This allows the data center to use data more efficiently and protect it over time. By automating these management tasks across a multi-tiered storage environment, they ensure that the right data is stored in the right tier at the right time to match the value of the data with the cost of the storage.

With turnkey applications to support email and DB2 environments and a wide variety of storage management solutions available, **STK has improved the performance, lowered the total cost of ownership, and simplified the administration of the storage system.**

Furthermore, they have enhanced storage automation with new hardware options for both disk and tape, in open systems and mainframe environments to facilitate the implementation of a multi-tiered storage architecture by adding low-cost disk arrays and fast access tape **to permit ILM applications to migrate data to the storage element whose cost best approximates the current value of the data.**



About The Clipper Group, Inc.

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

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

About the Author

David Reine is Director, Enterprise Systems for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 23. (Please dial "1-23" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

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

Disclosure

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

Regarding the Information in this Issue

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