by John Webster

September 10, 2004

Narrowing the Data Management Gap For Data-Intensive Applications

The data management gap—that is, the gap between how much data administrators needs to manage and the amount they can actually manage—is widening, especially among users of commercial applications such as OLTP and messaging.

Data volumes, meanwhile, are increasing across the board, and surveys show that, while hardware costs are decreasing, data management costs, on average, are rising between 25% and 30% per year.

Clearly, as storage environments become more complex, they become more costly and difficult to manage. Networked storage architectures (e.g., NAS, SAN, or both) are now commonplace and require a certain level of expertise to manage, which further adds to the overall data management expense.

While data in the average production IT environment is doubling every 2 years, in most data-intensive environments the amount of data that needs to be stored, managed, and made available is doubling every 6 to 18 months. To give an idea of the level of cost and complexity this can present, a data intensive environment that starts with 1 terabyte (TB) today will balloon to a petabyte in 5 years—the current life-span of typical storage hardware. And in some data-intensive industry segments such as pharmaceuticals, even greater challenges result from new state and federal regulations that dictate lengthy data retention periods of up to 30 years for certain data types.

Factor in the IT budget and the situation becomes even more daunting. Surveys show that while many CIOs were handed larger IT budgets for data-intensive computing applications this year, they were also asked to do more with proportionately less, which means that their workloads are, in many cases, still growing significantly faster than their supporting IT budgets.

www.datamobilitygroup.com

76 Northeastern Blvd. Suite 29A Nashua, NH 03062 Phone: 603.563.8384 Fax: 603.563.8282

To help rein-in some of these costs, many IT administrators are increasingly looking at new, relatively low-cost technologies, such as ATA/SATA disk and Linux. However, a looming question remains as to whether IT administrators in data-intensive environments will able to meet these challenges as they scale their storage environments to meet capacity demands going forward. Stated another way: How will these storage administrators narrow the gap between what they can and can't effectively manage with available resources in the coming year?

Parsing ILM and DLM

One way to narrow, if not eliminate, the data management gap is to manage information over its natural lifecycle. For the past year, systems and storage vendors alike have proposed this methodology, which is now generally referred to as information lifecycle management (ILM),

The idea behind ILM is that by matching data to the appropriate storage device in a tiered fashion, IT administrators will not only be able to manage more data more efficiently, but also increase the value of the data (hence IT) to their organizations. Application users benefit as well by getting their service level agreements with IT management met at a reduced cost because of the efficiencies that IT develops.

ILM has been proposed as a way to meet the major challenges forced on enterprises by regulatory compliance obligations. However, it is often the case that compliance-related data—email for example—can have little to no lasting value to the enterprise. The real value of ILM is missed if it is only viewed in this obligatory context. On the other hand, if ILM is seen as an efficient way to extract increasing value from all data objects as they pass through their "cradle-to-grave" lifecycles, then real ROI benefits resulting from ILM implementations can be seen.

But what is ILM exactly? Storage vendors typically define ILM as a collection of policies, processes, practices, services, and tools used to align the business value of information with the most appropriate and cost-effective infrastructure from the time information is created through its final disposition. Information is aligned with business requirements through management policies and service levels associated with applications, metadata and data.[1]In the context of ILM, that "cost-effective infrastructure" has come to mean a tiered storage hierarchy, consisting of primary (high-performance, high-availability disk), secondary (SATA disk and near-line tape libraries), and tertiary (low-performance, high-capacity tape or optical storage) storage systems.

But there are some problems with this definition. First, there is the issue of whether or not ILM actually exists today. After all, ILM is defined as a multi-layered process or architecture, not a single product. And then there is the fact that the concept, by focusing on "information" alone, ignores a very crucial difference between "data" and "information." We believe that this distinction is crucial in that it represents the dividing line between what can actually be implemented today (product-wise) and what is more philosophical or process-oriented.

76 Northeastern Blvd. Suite 29A Nashua, NH 03062 Phone: 603.563.8384 Fax: 603.563.8282

Before we look at what is "doable" today and what the future holds, we'll first address some fundamental questions like:

What is information?

What is data?

Is there a difference between the two?

What is data lifecycle management (DLM) and how does it differ from ILM?

Making the distinction

Computing has evolved from data processing (DP) to IT. However, somewhere along the evolutionary path from DP to IT, the words "data" and "information" have become near synonymous in many computing circles.

On the contrary, we believe these two words are distinctly different—particularly so in the context of data-intensive applications where saving data over extended lifecycles is required. Simply put, data is about bytes and files—it's the raw material. Information, on the other hand, is the refined product that results from the application processing the raw material (or data). To illustrate, note what happens to a Microsoft Word document over time. A document is created using a particular version of Word running a particular version of windows. Five years hence, the data that represents the document can still be accessed, but whether or not the document will be presented exactly in its original state is questionable unless one uses the same version of Word running under the same version of Windows to present the document to the user.

As such, ILM is a more complex process than data lifecycle management (DLM), and particularly so in data-intensive environments, requiring IT managers to "walk" massive amounts of data, along with the supporting systems, and applications through time in order to preserve information as it moves though its lifecycle.

DLM, in contrast, only requires IT managers to preserve the data throughout its lifecycle. However, we note that the need for scientific and medical communities to preserve research-related data for an "indefinite" time period is real and no less challenging.

Having made the distinction between data and information—in a sense deconstructing ILM—it is important to emphasize that, though distinctly different, data and information should be closely integrated over time and not managed as separate and distinct bubbles. In fact, we believe that DLM and ILM go hand-in-hand, with DLM as the first step of an ILM future.

Put simply, DLM is about data management services and currently available infrastructure components. These include:

- A storage infrastructure that consists of a logical continuum of hardware platforms, which runs from primary to secondary disk storage to long-term archival storage devices such as tape and optical media.

 76 Northeastern Blvd. Suite 29A Nashua, NH 03062 Phone: 603.563.8384 Fax: 603.563.8282

- Automated policy engines that are integrated with data movers and data copy functions such that data objects can automatically allocated to the most appropriate platform along the logical storage continuum based on pre-established policy.
- Data protection and risk-mitigation processes that address integrity issues during the entire lifecycle of data objects.

The following graphic illustrates the foregoing discussion:

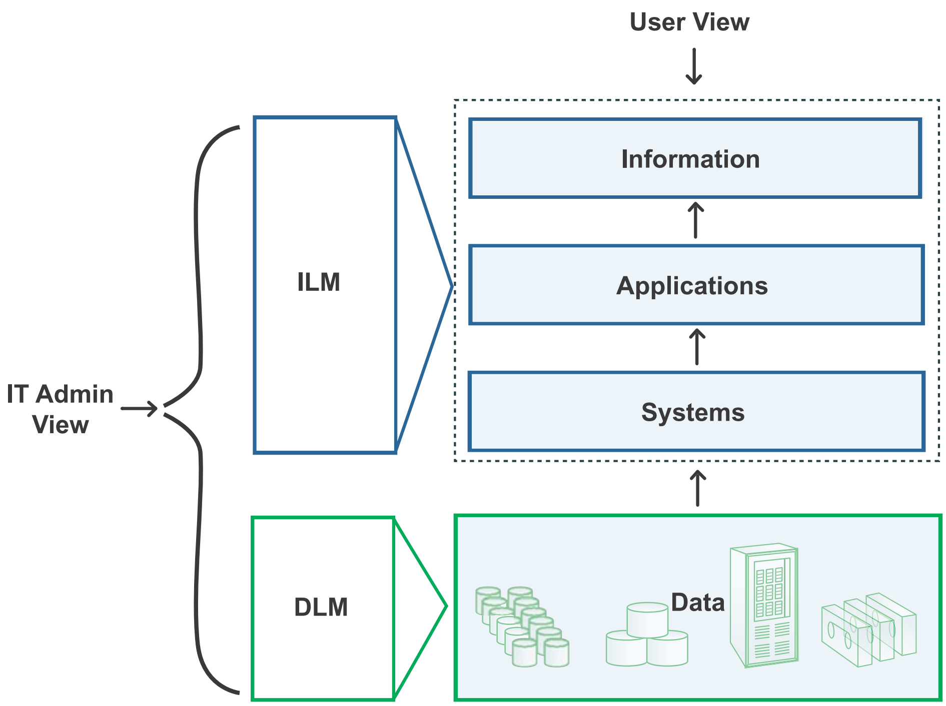

Here, the data intensive application user views information. Below that sits data, systems, and applications—the raw material and the processes that transform the data into information.

To the left and supporting that transformational process are a set of management services, which are seen by the IT administrators. DLM services relate directly to the data block and serve as a foundation for the ILM services, which relate directly to information.

ILM sits on top of and exploits the DLM data management processes. DLM processes move, replicate, reorganize, and archive bits, bytes, and blocks of data. In essence, DLM integrates with ILM to provide essential services to ILM.

The logical storage continuum

We believe that, for data-intensive applications, data center managers can lower their overall storage (hardware plus software) and storage management costs from 60% to 80% by simply moving static and less-frequently-accessed

76 Northeastern Blvd. Suite 29A Nashua, NH 03062 Phone: 603.563.8384 Fax: 603.563.8282

data off primary data storage and onto more appropriate secondary or tertiary storage tiers. This can be done transparently, via automated or manual processes, without disrupting users or applications.

As data ages, users typically access it less and less frequently and its relative value to an organization generally decreases. Therefore, it makes little sense to keep low-demand files with probably less organizational or research value on the same high-performance storage systems as high-demand files.

It is estimated that for every 100 TB of total storage (disk plus tape) in data intensive environments, 2TB to 10TB of data resides on primary disk (with the average being 3 TB)[2]. The bad news is that when administrators audit the access patterns of data objects residing on high-performance disk, they typically find that 80% of I/Os are hitting about 20% of the data (the 80:20 rule). This means that at any given point in time, the bulk of the data residing on a high-end disk system (approximately 2.4 TB) is essentially static.

The good news is that existing storage hardware platforms used for backup and restoration can also serve as the foundation for a storage continuum where static data and less-frequently-accessed data objects are transparently moved back and forth from primary storage (high-performance disk) to secondary (SATA disk and nearline tape) to tertiary (lower performance tape and optical) storage tiers as needed.

Data movers and data copy functions

DLM assumes that a process is in place that can move data objects bi-directionally along the storage continuum as required. This can be done in two ways: using copy functions that create multiple copies of data objects, and using data movers that migrate data objects from one device to another device along the continuum. Copy functions create a duplicate (or multiple) copy(ies) of data to be stored on the appropriate storage tier and LM services can then concurrently track both primary data objects and all associated copies. Data movers, on the other hand, relocate data objects among the storage tiers without creating duplicate copies. A file created on one device is moved or migrated to another device as the data ages and its business relevance changes. We believe that, in data-intensive environments, both functions are required and can be automated or initiated manually. Copy functions provide copies for data protection, while migration functions move data among resources in the continuum to the most efficient storage tier.

Automated policy engines

Automated policy engines harness the power of data movers. In essence, they tell the data movers when and where to put the data. Minimal or no human intervention is needed once IT administrators have categorized and set policies for the data. They are transparent to the user application environment and they free up valuable staff time and reduce data loss due to human error.

76 Northeastern Blvd. Suite 29A Nashua, NH 03062 Phone: 603.563.8384 Fax: 603.563.8282

for example:

SGI InfiniteStorage Data Lifecycle Management Server

Silicon Graphics, Inc. (SGI), a well-established vendor in HPC environments, recently introduced its pre-configured InfiniteStorage Data Lifecycle Management Server.

The solution combines SGI Data Migration Facility software and InfiniteStorage SAN or NAS servers. These servers are based on Intel 64-bit Itanium 2 processors for high-performance metadata processing—an important attribute for sustaining high data rates while supporting data-intensive applications. As disk subsystems double in capacity every 12 months, the cache must keep pace, and the high-end systems will require 64-bit addressing. The Intel Itanium processor family will enable these kinds of scalable systems such as the SGI MetaData Server

Optionally, the server can be configured with SGI Cluster XFS (CXFS) for DLM capability in heterogeneous application server (e.g., SGI IRIX, Linux, IBM AIX, Sun Solaris, and Microsoft Windows) environments.

With this product, file migration occurs in two stages:

> Stage One: A file is copied (migrated) to secondary storage either tape and/or lower cost SATA RAID.
>
> Stage Two: After the copy is secure, the file data blocks are released on the primary storage device and are returned to the free space pool. Files are deemed fully migrated when all near-line copies are complete.

The DLM Server supports multi-tiered, hierarchical storage management (HSM) processes that use high-performance primary disk, lower-cost near-line disk (SATA), and tape. Data moves automatically and transparently up or down the storage hierarchy based on user-defined criteria, or standard policies such as time of access, size, owner, and group.

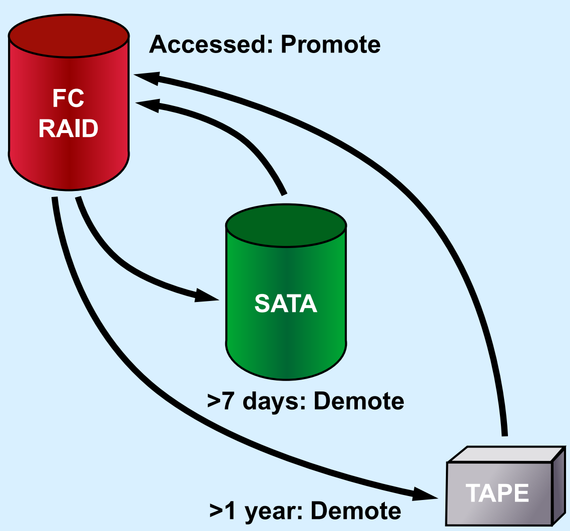

76 Northeastern Blvd. Suite 29A Nashua, NH 03062 Phone: 603.563.8384 Fax: 603.563.8282

The software also interoperates with a variety of third-party high-availability, backup, and shared file-system solutions. DMF automatically manages the free space on primary disk storage to ensure that disk space is always available.

Any user or application request to read or write data in a file system managed by the DLM Server may trigger a recall of a file to the primary disk if that file has been migrated. File recall occurs automatically if a file has been migrated.

In the initial release of the product, a user's recall request is delayed while the file is recalled to disk, otherwise all activity is transparent. The next release of the product (due out later this year) will allow users to recall files without affecting the performance of the user application.

Additional noteworthy features of the DLM Server include:

Consistency and integrity—DMF maintains file-system metadata databases to record all information about stored data that are synchronized with the file systems that the server manages. DMF can be configured to automatically examine the consistency and integrity of these databases. Periodic backup of these databases is performed or automatically via CLI scripting. Administrtor's can also perform manual backups to supplement high activity periods.

Multiple media support—Near-line media is the destination of all data migrated from primary disk and is managed by daemon-like DMF components called the mediaspecific process (MSP) and the library server (LS). Multiple copies of each file can be created (two copies is standard but up to 64 separate copies are supported) and these can be directed to different MSPs. MSPs are available for:

- Disk—Provides migration to secondary disks (Fibre, SATA, or other) and other random access devices like NFS-mounted file systems
- FTP—Supports migration of data across IP networks to local or remote locations
- Tape—Tape libraries are the most common migration target currently in use

Simplification of backup processes—The DLM Server moves only data associated with files, not file indices or directories. The backup process can be simplified by creating a policy of migrating 100% of DMF-managed file systems, leaving only a small amount of data to be backed up. Backups can be performed using the IRIX xfsdump and xfsrestore utilities (DMF includes tools to simplify and automate execution of these utilities) or other backup programs that are DMF and DMAPI- aware.

Parallelism—Multiple operations occur in parallel so that user processes aren't impacted while other operations complete. Maximum migration bandwidth is typically only limited by the aggregate bandwidth of attached near-line storage.

76 Northeastern Blvd. Suite 29A Nashua, NH 03062 Phone: 603.563.8384 Fax: 603.563.8282

Scalability—The DLM Server server architecture can manage up to 2 billion migrated files at one time, 18 million terabytes (18 exabytes) in a single instance, and single file sizes of up to 9 million terabytes (9 exabytes).

Before installation, data that will be managed by the DLM Server is categorized based upon the user's expectations for length of retention and swiftness of recall, and any chargeback agreements that may apply. The Administrative staff then establishes policies for managing data on a per-category basis, perhaps using different policies for different projects. The DLM server is then configured to implement these policies. Initial configuration also requires IT administrators to choose the appropriate hardware for the storage tiers (i.e., primary disk, secondary disk, and tape (volume groups offering different characteristics for capacity, streaming and time to first byte).

Attention should also be paid to the expected rate of recalls by users. This assumption is used to design the requried bandwidth between the DLM server and the underlying storage subsystems to ensure that data can be migrated and recalled at a rate necessary to achieve operational goals. SGI Professional Services are available to help customers tailor the DLM Server environment to their specific requirements.

Recommendations for users

What is the value of DLM to HPC IT administrators? The bottom line is that DLM can help IT administrators today make better use of their high-value disk assets by moving data that is of declining relevance and business value to secondary or tertiary data tiers along the storage continuum.

DMG recommends that HPC data center managers carefully consider the following issues when evaluating a DLM implementation project.

Can you automate and thereby eliminate manual data recovery processes (typically from tape) that are almost always more costly on a TCO basis than automated processes? In addition, how many other manual processes can be automated?

Can you further reduce both planned and unplanned downtime using DLM services to provide continuous data accessibility and continuous data protection?

Can you reduce the amount of data residing on primary storage that requires the creation of backup copies?

Can you reduce the response time when recalling data objects from tape?

76 Northeastern Blvd. Suite 29A Nashua, NH 03062 Phone: 603.563.8384 Fax: 603.563.8282

Does the proposed DLM solution help you more accurately predict true storage capacity requirements based on actual asset usage, and do so before you run out of free space? Can you add capacity in tiered increments on an as-needed basis - and at a cost that is reflective of current market prices? And does the solution scale to the petabyte range sometimes found in data-intensive environments?

Is data integrity assured via functions such as media checking, recovery processes, journaling, and implementation of two-phase commit processes?

Can it support new secondary disk storage technologies such as ATA disk, thereby creating multiple levels of data protection, in the context of both backup and restore, and DLM?

Conclusion

DLM provides management services that allow HPC IT administrators to make more efficient use of their high-value disk storage by migrating data that is of declining relevance and business value to secondary storage or tertiary storage tiers.

Through policy-based automation, organizations can also make better use of their IT staff and minimize application downtime by providing continuous data accessibility and continuous data protection. DLM-automated processes can also eliminate manual data recovery processes that tend to be more costly on a TCO basis.

Finally, DLM services can help IT administrators more accurately predict true storage capacity requirements based on actual asset usage. Capacity can be added in tiered increments on an as-needed basis—and at a cost that is reflective of current market prices.

Footnot s

[1] From SNIA ILM TWG definitions document dated 7/28/04. We applaud the SNIA's effort to establish standard ILM-related definitions as the topic is complex and the terminology is open to misuse by vendors.

[2] Source: SGI customer survey data.

 76 Northeastern Blvd. Suite 29A Nashua, NH 03062 Phone: 603.563.8384 Fax: 603.563.8282