

Understanding How ILM is Evolving



THE INFO PRO
www.theinfopro.net



This TheInfoPro (TIP) Research Paper delivers findings from one-on-one interviews with Fortune 1000 storage professionals. Every six months TIP releases a Storage report based on over 150 interviews with Fortune 1000 enterprises. Most of the data used in this paper was collected for Wave 6 in Fall 2005 and Wave 7 in Winter 2006.

Introduction

Compellent's Storage Center is the **industry's ONLY SAN with Automated Tiered Storage**, eliminating the manual data classification and manual movement of data between tiers associated with other ILM solutions.

Storage Center's Data Progression is a complete hardware and software solution that fully automates tiered storage.

Data Progression automatically classifies and migrates data to the optimum tier of storage based on usage allowing businesses to retain frequently accessed data on high performance storage and storing infrequently accessed data on lower cost storage.

With Data Progression, companies are able to:

- Reduce their disk expenditures by as much as 74%
- Cut storage administrative time in half
- Improve business continuity by storing up to 8 times more Replays (snapshots) at no additional cost
- Increase performance by moving inactive data off Fibre Channel drives.

ILM as a Growth Management Strategy

TheInfoPro has been interviewing storage professionals for over four years, and in the last 18 months interest in ILM has heated up. The difficulties of managing growth are becoming more prevalent of late as Storage organizations struggle to meet ever increasing demands for faster application performance, longer regulatory compliance retention, and shorter backup windows. The industry is responding with increased interest in ILM as well as those technologies that complement it, such as data classification, categorization, and virtualization, and embedded data migration.

Storage Management Technology Heat Index™ Rankings


Technology	Spring 2005	Fall 2005	Spring 2006
Information Lifecycle Management (ILM)	4	1	1
Online Data Migration Software / Data Mobility Software	11	4	2
Data Classification / Indexing / Categorization Software	N/A	6	3
Data Encryption Software	N/A	14	3
Email Management and Archiving Software (File)	1	3	5
Network File Virtualization (Unstructured Data)	N/A	15	6
Storage Tiering Software (Automated Storage Provisioning)	7	11	6

The TIP Heat Index™ is based on the immediacy of user need – their plans for each technology – and weighted by their storage spending. ILM is rated as the hottest technology.

Storage Capacity Is a Train That Has No Stops

Consolidation projects are topping the list of priorities for Fortune 1000 storage professionals in 2006. The focus is the result of a storage binge that took place in 2005 where the average usable SAN footprint among Fortune 1000 companies doubled, nearing 500 TB. With this corresponding growth, storage pros have been looking for technologies that can help them rationalize their installed storage and help them end-of-life older, higher maintenance legacy equipment while at the same time meet the increasingly stringent storage standards.

Usable SAN Capacity in Terabytes



A storage binge that took place in 2005 where the average usable SAN footprint among Fortune 1000 companies doubled, nearing 500 TB.

Comments from F1000 Storage Pros:

- "The first issue for us is around the fact that our applications continue to grow while our recovery time objectives continue to be the same. We are having difficulty getting the data back to them in this same RTO timeframe. The other is around the management of the storage, which continues to be a challenge given the size of our storage." (Telecom & Technology)
- "The key dynamic causing growth in Networked Storage is the influx of Sarbanes Oxley and HIPAA. We are also in a government compliance program which necessitates doing a massive amount of reporting and data retention." (Healthcare)

How It Works

Automated Tiered Storage does not require an administrator for data movement or monitoring, and, therefore, is able to significantly reduce storage administration time by eliminating data classification, complex software requirements and all integration issues.

And, because Data Progression is fully integrated with Storage Center, there is no need to purchase, install, train and learn costly third-party hardware and software products.

Data Progression utilizes meta-data collected by the SAN to automate data classification and intelligently migrate the inactive portions of large files to lower-cost disk.

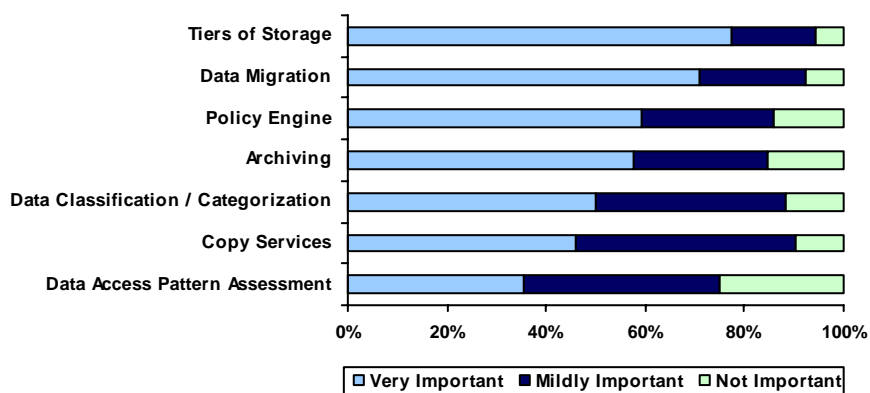
Unlike traditional storage systems that move entire files or volumes, metadata enables Data Progression to intelligently migrate data at the block-level for optimum results.

Data Progression eliminates the need for ongoing, time-intensive, company-wide data classification that impedes ILM strategies and negates any possible return on investment.

Phases of ILM Deployment

When Storage professionals were asked in 2005, what the most important components of a successful ILM strategy were, the following priorities made the list: 1) Identification of how their applications are using storage, and create a tiered storage architecture that is optimized for the changing value of data in the enterprise; 2) Finding data migration tools to help move aged or higher throughput required applications between the tiers; 3) Codification of the rules used for these migrations into the standard operating procedures and policies of the organization; 4) Linking the backup and archiving components together into the overall ILM framework.

Importance of ILM Functionalities



Above is a list of priorities from fall 2005 when storage professionals were asked to name the most important components of a successful ILM strategy.

ILM Functionalities Prioritization

While storage professionals were struggling with data classification and data categorization, email professionals were struggling with classifying and retrieving email for regulatory compliance and in particular using Content Addressed Storage as way of addressing the meta indexing and searching requirements. It didn't take long for the two groups to realize the importance of data classification for ILM holistically and by early 2006, data classification was listed as the most important functionality of an ILM Strategy.

Comments from F1000 Storage Pros:

- "Seamless data migration capability is the most important functionality [for ILM] and it is only borderline available today." (Financial Services)
- "We need to be able to determine where or what you have stored and report on it. I have a lot of concerns about how ILM operates in an Open Systems arena." (Industrial & Manufacturing)
- "My biggest pain point is getting a handle on how much data I have and where it is, and also relating that data back to the business." (Other) "The business understanding why they need to have all the different types of storage is my biggest pain. For ILM, tiered storage is just about moving data. The business doesn't really care. It's not a benefit or challenge. I can't explain technology to the business side." (Financial Services)

Key Features

Compellent's **Dynamic Block Architecture** records and tracks specific information called "metadata" about blocks of data, enabling intelligent data movement.

Automated data classification and block-level migration obsoletes manual classification and migration, saving valuable administrative time.

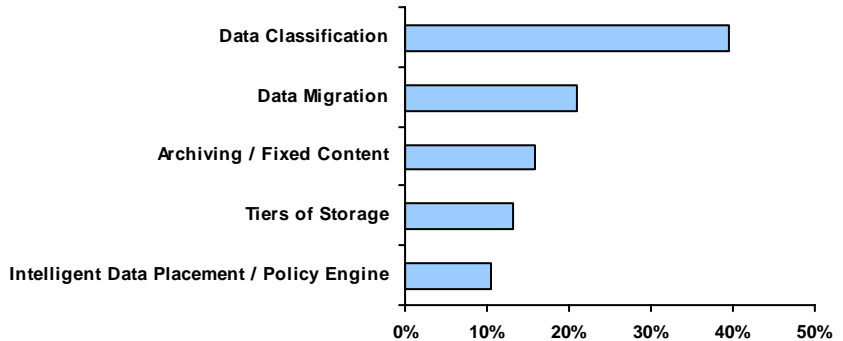
With Compellent's **standards-based, modular architecture**, storage tiers can utilize any technology and scale to any capacity. This **technology independence** allows you to mix and match disk drives, using the right blend for each tier.

Compatible with any OS, drive type or application, Data Progression even works seamlessly with large databases.

The **intuitive wizard-base interface** allows rapid set-up and management, making it easy to understand where data is stored and determine ongoing disk requirements over time.

Data Progression **automatically migrates space-efficient Replays (snapshots) to a less costly class of storage** maintaining immediate availability for reliable recovery and the ability to store 8x more Replays at no additional cost.

Most Important ILM Functionality

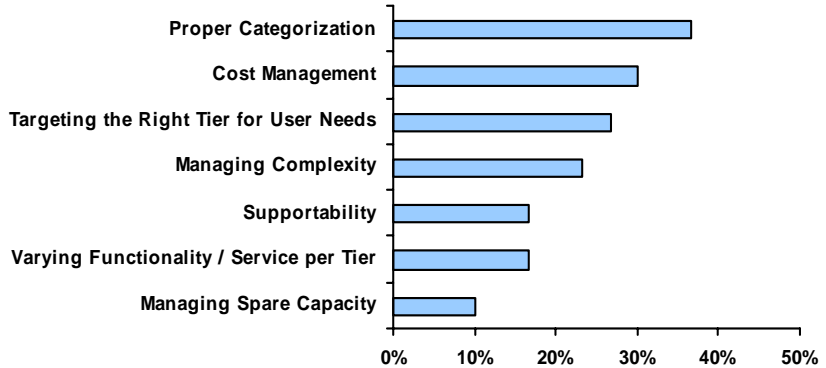


Data classification was listed as the most important functionality of an ILM Strategy in our Fall 2005 Storage study.

Data Mobility Pains

While ILM strategies appeared straightforward, storage professionals were struggling to execute data migrations. Application changes and marketing dynamics were continuously alternating the valuation of data, and migration activities were manual and required intense coordination with application teams to execute transparently. Storage professionals needed a more intelligent way of discovering and moving data transparently.

Biggest Challenge for Tiered Storage



Storage pros struggle to move data between tiers because of an inability to intelligently discover and define data first.

Comments from F1000 Storage Pros:

- "Archiving and managing storage growth are interrelated in our organization. We are going through a transformation right now and we've been looking at optimizing our storage environment and we see data classification as the missing link."
- "Classification is very user driven, based on personal preference. We have better technology but implementation is still based on getting the right information. Users are not qualified to give criteria and storage staff doesn't have the authority." (Telecom & Technology)
- "We have implemented tiers of storage, but not data classification. We don't have an intelligent data placement engine. We do some business archiving, but not from a systematic unstructured perspective. I don't believe that engine exists at this time." (Industrial & Manufacturing)

Customer Results

“Data Progression has moved two-thirds of our inactive data – both documents and e-mail – onto a less expensive tier of storage automatically with no impact to the end-user.”

Jeff Barnes, Briggs and Morgan

“It is impossible for any organization to manually track and determine when any single piece of data has grown inactive. Data Progression automatically migrated 85% of our data to lower-cost SATA drives without any manual intervention, resulting in an immediate payback and increased performance and availability on our Fibre Channel disk.”

Ryan Fletcher, Whitehall City Schools

“With Data Progression, we have cut our video production time by 50%, saved 75% on storage costs and increased capacity by adding a third tier of SATA storage that serves as online archival backup.”

Kevin Pazerra, Maine Public Broadcasting

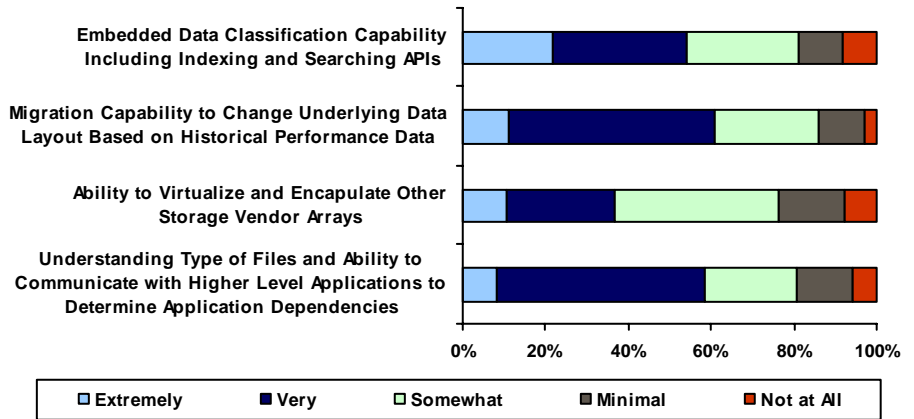
“I cannot imagine trying to do ILM without a tool like Data Progression, it would mean spending hours a week looking at data usage, deciding what to move and then taking the time to move it. Data Progression saves us many, many hours in administrative time.”

Rodger Smith, IWCO Direct

Next Generation Storage Subsystems

By the Spring of 2006, data classification needs had started to become the most common theme for next generation storage equipment. The desire to link application valuation with storage system allocation was listed as the top priority for Tier 1 storage subsystems.

Importance of ILM Functionalities in Future SAN Storage Array Products



The desire to link application valuation with storage system allocation was listed as the top priority for Tier 1 storage subsystems in our most recent Storage study, completed April 2006.

Conclusion

As Fortune 1000 businesses continue to use information to differentiate themselves, deployed storage will invariably continue to grow, and with that growth comes increased complexity, physical space, and higher costs. In the short term, Storage organizations are using data mobility and trying to end of life older, higher maintenance equipment to help alleviate a budget that is already under intense pressure. Unfortunately, trying to determine which components are underperforming and should be targeted, is not an easy task. According to Wave 7, there are a handful of ILM functionalities that are starting to gain interest to help with this problem. Data classification, categorization and indexing are starting to emerge as critical components of a successful ILM strategy and storage professionals are indicating a desire to see these features embedded in next generation storage array products. Whether these next generation products will finally close the gap between storage system data silos and business value is still to be determined but what is clear is that storage professionals are becoming more interwoven into business management and the competitive differentiation of data usage.