



Comprehensive ILM Initiative Delivers Business-Oriented Information Infrastructure for EMC

Table of Contents

Comprehensive ILM Initiative Delivers Business-Oriented Information Infrastructure for EMC	3
Efficiency and cost challenges greet new CIO	3
IT turns to ILM for answers	3
Phase I: Efforts Focus on Building a Tiered Infrastructure	4
Classification aligns storage with business need	4
Simplified approach yields clear, objective decisions	4
Consolidation delivers scalable, efficient, and more manageable infrastructure	5
Tiered service catalog eases consolidation moves	5
EMC ControlCenter allows IT staff to manage consolidated infrastructure in real time	6
More reliable backup/recovery procedures lower risks, improve performance	6
CLARiiON disk libraries gradually replace tape backup	6
One-hundred percent reliability on restore requests achieved	7
Greater awareness of business continuity cost/benefit tradeoffs leads to more objective decisions	7
Phase II: Applying ILM to Key Application Environments	8
ILM improves performance and lowers operating costs for key Oracle applications	8
Making sure users understand archived information is still accessible	8
Users much happier, business more efficient after implementing automated archive	9
ILM facilitates Exchange migration, helps standardize e-mail policies	9
Archive delivers virtually unlimited mailboxes from user's perspective	9
Phase III: Using ILM to Simplify Information Delivery	10
Unstructured documents increasingly important for critical business processes	10
Automated policies help ensure document currency and accuracy	11
“Embedded” ILM Has Changed the Conversation between IT and the Business	11

Comprehensive ILM Initiative Delivers Business-Oriented Information Infrastructure for EMC

Efficiency and cost challenges greet new CIO

When a new CIO arrived at EMC Corporation four years ago he was surprised at how difficult it was to get consistent performance and utilization statistics from the five primary data centers. What he experienced was the downside of EMC's recent successes. EMC had grown substantially over the past few years including several recent acquisitions. Unfortunately, ad-hoc growth of its information technology had led to a complex, disparate high-end storage infrastructure that was difficult to manage, expensive to operate, and now had a limited ability to scale to meet growing demands.

IT turns to ILM for answers

The IT executive team concluded that deliberate action was needed to put more effective architectural and operational standards in place. In order to reduce costs and improve the efficiency of storage assets, EMC required a services-driven infrastructure and IT organization across EMC's five primary data centers. Information lifecycle management (ILM), which was still an emerging strategy in 2003, provided the framework for what proved to be a comprehensive initiative, involving many parallel, related work efforts. As one of the original proponents of ILM, it was a natural step for EMC to begin the ILM journey and practice what it had begun to vigorously preach.

The overall objective of the ILM Showcase, as the multi-phased project came to be known, was to better align critical storage resources with the needs of the business and improve efficiencies, including reducing and avoiding costs to deliver information wherever possible. The ILM Showcase project covered three phases and addressed seven core IT disciplines in total. The first phase focused on the development of a tiered information infrastructure and involved **classifying** and **consolidating** critical information resources; improving the reliability of the **backup, recovery, and archive** processes; and evaluating **business continuity** capabilities. The second phase focused on more detailed classification and policy development, extending ILM across **key application areas including Oracle and Exchange** environments. The third phase involved the implementation of a **content management** infrastructure that focused on bringing EMC's unstructured content into the context of ILM. The objective of this phase was to integrate this unstructured data with the structured content of the application environments to better support key business processes, such as customer support and employee communications.

① Phase 1: Tiered Infrastructure	Classification, Consolidation Backup, Recovery, Archive Business Continuity
② Phase 2: Application- specific ILM	Oracle, Exchange
③ Phase 3: Cross-application ILM	Content Management across applications

Figure 1. ILM Showcase Project: Three phases

EMC® Consulting helped guide the eighteen person project team—over half of which came from the IT organization—in the effort. The scope of the ILM Showcase was ambitious by most measures. With nearly one petabyte of storage, this was easily one of the largest efforts for the consulting group when compared to its most significant customer engagements. As an enterprise-wide initiative, it provided an opportunity to apply a comprehensive

ILM strategy across EMC's entire information infrastructure. This paper chronicles the project team's experiences over the last three years and explores the learning that took place. While many events did not happen exactly as planned, the ILM Showcase illustrated the differences that occur between an ideal environment and reality and ultimately provided some very important and affirming lessons about the value of ILM.

Phase I: Efforts Focus on Building a Tiered Infrastructure

Classification aligns storage with business need

One of the primary IT issues EMC faced was the high average cost per gigabyte of its storage infrastructure. The IT organization did not have a process for evaluating business needs and then mapping these to appropriate storage services. So, the storage infrastructure did not effectively reflect the requirements of the business processes it was designed to support. The result was that a majority of the capacity was high-end storage even though midrange or lower platforms were adequate for supporting the business requirements in many cases. By classifying key information assets and applications EMC would be able to meet the storage requirements of key business processes far more cost effectively at a much lower average cost per gigabyte. In the process, less overhead-heavy RAID technologies would be utilized which would further increase the utilization of raw storage capacity.

Two years ago when the classification effort began, there was a smaller body of intellectual property surrounding ILM. As a result, the IT organization and EMC Consulting had to educate key application and business constituents during the process. The most significant challenge in starting the classification effort was convincing the user community that the benefits outweighed the effort required to go through the applications and data to evaluate the business service requirements for each. For many classification initiatives, it is common for the user community to feel they cannot afford the cycles or resources to go data set by data set and have difficulty understanding the value of the exercise. When this happens, the consulting resource and IT must negotiate with application owners to determine how to attack this challenge.

Simplified approach yields clear, objective decisions

By dividing the problem into small steps and tackling one at a time, the IT organization and EMC Consulting were able to experience success across a huge pool of data. One of the first decisions made was to take a more macro approach by looking at applications and evaluating the service levels required. It is a more fundamental analysis that makes tiering decisions based on the overall application service level and drags all other application components with it. As it turned out, this approach proved invaluable in the understanding IT gained regarding the business application's service requirements. By taking technology and cost considerations out of these discussions, for example, focus shifted to the business process that was being automated by the application and the business impact of IT service disruptions. The result was a more comfortable conversation for the business owners that was more objective and devoid of the kind of emotional decisions that usually take place. By adding technology and cost considerations back in after business impacts were determined, their requirements were mapped back to infrastructure tiers. Business users were guided to make very clear, unemotional decisions that made sound business sense.

“There is a point of diminishing returns in deciding how detailed to make the classification analysis. Often the more simplistic approach yields 80 percent of the benefit and avoids the complexity that scares many business users away from classification projects before they can get started. There is no ‘easy button’ so it is important to start at the application layer and drill deeper from there as it makes sense.”

— Jeff Gabriel, EMC Consultant

“There is a point of diminishing returns in deciding how detailed to make the classification analysis,” explained Jeff Gabriel, the lead EMC Consultant on the project. “Often the more simplistic approach yields 80 percent of the benefit and avoids the complexity that scares many business users away from classification projects before they can get started. There is no ‘easy button’ so it is important to start at the application layer and drill deeper from there as it makes sense.”

As a result of the classification analysis, many applications were shown to be receiving service levels higher than required and a significant percentage of data across the enterprise resided on storage tiers that provided higher service levels than necessary. Significant savings were achieved by moving these applications and data to lower service-level tiers. Application alignment and classification ultimately drove EMC to a more heterogeneous/blended environment. The implementation of EMC's automated storage software technologies allows the IT organization to operate this physically more complex environment with greater movement of data because it is now more cost-effective and flexible from an operational standpoint.

“There is always a balance between greater complexities and the benefits derived from cost savings and greater flexibility. We had issues in getting the user community on board with the classification effort, but we learned a great deal in how to attack the challenge. How do we overcome the resistance to change? How do we identify the applications to target? Which applications should be set aside? We discovered the answers to these questions the tough way and EMC is now able to leverage this experience for its customers’ ILM efforts.”

— Dan Inbar, EMC Senior Director of Global IT Technology

“There is always a balance between greater complexities and the benefits derived from cost savings and greater flexibility,” explained Dan Inbar, EMC’s Senior Director of Global IT Technology. “We had issues in getting the user community on board with the classification effort, but we learned a great deal in how to attack the challenge. How do we overcome the resistance to change? How do we identify the applications to target? Which applications should be set aside? We discovered the answers to these questions the tough way and EMC is now able to leverage this experience for its customers’ ILM efforts.”

Consolidation delivers scalable, efficient, and more manageable infrastructure

With a better understanding of the true business value and impact of critical information assets and the storage platforms that should support them, the focus turned to consolidating the widely scattered storage capacity. The current ad-hoc infrastructure had significant operational issues and could no longer scale to meet the data growth projections. The lack of an enterprise view of storage limited the ability to make effective decisions regarding storage and often led to redundant and unnecessary capacity acquisitions and low overall capacity utilization. Complicating matters further were the myriad of inefficient and error-prone manual tools and processes that were in place to support storage, which resulted in unsynchronized processes and unclear responsibilities across infrastructure groups. This confusion led to increased costs, reduced quality of service, and added to the scalability issues.

Having the ability to be more proactive on failure incidents — the sheer volume of drives installed presented EMC with a statistically unavoidable loss of 10-12 drives per week — was another critical objective that could be addressed by upgrading to the latest version of EMC ControlCenter® software. The scale of the consolidation effort was perhaps the biggest challenge. “We took on 100 percent of the environment,” said Bill Pratt, EMC’s Director of Enterprise Storage. “Most customers would not have attempted this amount of infrastructure, so we were navigating uncharted waters.”

Tiered service catalog eases consolidation moves

Because the classification effort had made the technology decision process more objective for business users, IT was able to avoid many of the technology biases, such as the preference to store Symmetrix® business unit data on Symmetrix storage and CLARiiON® research and development data on CLARiiON systems, which would have otherwise threatened the consolidation effort. Whenever IT recommended moving data to a lower tier of storage as part of the consolidation process, the service-level catalogue clearly showed the relationship between the business units’ requirements and the costs to deliver an appropriate service level using the lower storage tier.

With the service-level catalogue in place, the applications were more easily aligned with the infrastructure tiers. Outside of the approximately 200 applications in the first tier, there were more than 250 additional mid-tier applications delivered by over 2,000 Windows and 1,000 UNIX servers to be considered in the consolidation.

“It was a mutual education for our IT organization and the consulting group,” explained Pratt. “Consulting was confronted with some challenges unique to EMC IT, as an internal customer. Whenever the theoretical underpinnings of the recommendations didn’t match our operational reality, they were able to adapt their approach to accommodate these challenges, particularly in the application alignment and consolidation process. In the end, EMC Consulting provided us the methodology, tools, and program management we needed to complete the effort.”

EMC ControlCenter allows IT staff to manage consolidated infrastructure

IT took longer than anticipated to complete the consolidation effort due to a variety of factors such as the limited availability of windows of time to move data. Nevertheless, the results have been substantial. Consolidation and realignment has moved many information assets from high-end storage units to more cost-effective mid and lower tiers. IT went from supporting 250 applications, all with different service-level agreements (SLAs) to a very manageable five committed service levels. SAN fabrics went from 63 down to 13 and high-end storage frames were reduced significantly from 115. NAS cabinets, which at 150 TB account for ten to fifteen percent of the SAN capacity, went down as well with a four-fold reduction in data movers and storage volumes.

In addition, the incorporation of the latest EMC ControlCenter technology has substantially improved storage management by significantly reducing the overall number of management points and lowering unit costs. The new architecture has quickly proven to be scalable and flexible, satisfying multiple storage demands in support of new enterprise projects. New remote management capabilities allow IT staff to know in real time when a drive has gone down. Now, service actions can be taken immediately rather than hearing from a very frustrated user later. And while all of the provisioning objectives have not yet been achieved, it has been a huge gain from a reporting standpoint. “We are getting very rich performance data that we didn’t have before,” stated Pratt. “Complete reports can now be generated in only fifteen minutes. The consolidation effort coupled with the ControlCenter implementation has allowed us to do much more with fewer resources.”

“We are getting very rich performance data that we didn’t have before. Complete reports can now be generated in only fifteen minutes. The consolidation effort coupled with the ControlCenter implementation has allowed us to do much more with fewer resources.”

— Bill Pratt, EMC Director of Enterprise Storage

More reliable backup/recovery procedures lower risks, improve performance

EMC Consulting and IT also addressed backup, recovery, and archive challenges in parallel with the consolidation and classification effort. Over the years, EMC had implemented many different backup mechanisms, including Veritas and Legato. Not surprisingly, the backup procedures in place were inconsistent. The IT organization’s most significant challenge was the limited reliability of the tape backup processes. Any concerns IT had about backup windows were dwarfed by the integrity of the backup tapes.

“We always seemed to be in triage mode,” explained Pratt. “More than 30 percent of the requests to restore files couldn’t be fulfilled because the backups had ultimately failed.” Although archive procedures should normally be addressed first to reduce backup dataset volumes and meet backup window constraints, Pratt just didn’t feel comfortable tackling archive while tape backup processes were failing. “Survivability is the primary concern for any IT organization,” continued Pratt. “Our objective, first and foremost, was to develop a consistent and reliable enterprise backup strategy.”

CLARiiON disk libraries gradually replace tape backup

In the fall of 2003, when this project began, moving 250 TB of production data over the network using remote disk backup was an intimidating proposition for the IT team because backup-to-disk technologies were just being introduced. Since then, the backup and recovery effort has evolved as new products have become available. For EMC, the availability of the CLARiiON Disk Library (CDL) was the turning point for the feasibility of backup-to-disk. CDLs have gradually replaced aging tape libraries for transaction database applications including the CRM, ERP, customer support, decision support, and ad-hoc reporting systems.

For other applications, the combination of lower cost fixed-content storage and sophisticated archive software has proven invaluable. One critical manufacturing application controlling production in the Cork, Ireland facility was generating progressively more detailed manufacturing test data that had to be stored on a historical basis. The increased data volumes had begun affecting the performance of the production systems

and the manufacturing process itself. In order to address this issue, policies were developed to integrate an active archive capability with the backup and recovery procedures. EMC DiskXtender® software was set up to automatically move historical production data more than 90 days old, which turned out to be more than 75 percent of the Cork production data, onto a local Centera™ system. This data was then asynchronously replicated to EMC's manufacturing site in Franklin, MA for added protection.

One-hundred percent reliability on restore requests achieved

The benefits that resulted from the more reliable disk-based backup processes extended well beyond the obvious reduction in number of backup tapes needed. For one mission-critical application, the time to execute a full production dataset backup went from seven hours to one. There were also reductions in required disk for many applications because a business continuance volume was no longer needed. The time required to recover files has improved from 9 GB an hour to 45 GB an hour. Where CDLs were installed, the benefits were even more significant because the implementation had virtually no effect on the backup process so no training was required.

The transition from tape to disk backup has been an unmitigated success from a reliability standpoint. "Our priority was to get a stable backup infrastructure in place and we accomplished that with flying colors," stated Pratt. "Now virtually 100 percent of the requests for file recovery can be fulfilled and we've experienced significant performance improvements in many mission-critical applications."

Greater awareness of business continuity cost/benefit tradeoffs leads to more objective decisions

EMC had actually begun to address business continuity, the last element of the first ILM phase, in the fall of 2001 before the ILM Showcase was formally initiated. The IT organization subsequently incorporated the existing business continuity initiative into the project. Shortly after the events of September 11, 2001, EMC was unsure of its information protection levels and felt compelled to reassess its business continuity plans. For EMC, the objective of the assessment was to educate business units on actual protection levels. In an ideal ILM implementation, classification of applications would take place first in order to identify business value, which would determine the relative levels of protection required. In this case, there were 200 known applications to be evaluated across functional groups. Without the luxury of classification, a decision was made to focus on the 80 to 100 applications that covered critical order-to-cash processes. Of these, 15 to 20 were determined to be the most mission-critical applications.

The primary example was the Oracle 11.03 ERP system that controlled applications supporting the order and financial management processes. With many legacy application interfaces, it was a heavily customized environment that experienced significant utilization spikes toward the end of each quarter so the business impact of a disruption would fluctuate over time. New regulatory requirements brought on by Sarbanes-Oxley had raised the stakes even further for the information assets managed by this complex system. As might be expected from a storage company, there had been a greater focus on data loss than recovery windows so very little data would be lost in the event of a disruption. Unfortunately, evaluation revealed that the time to recover these applications was unacceptably long.

The results of the assessment were eye-opening for many of the business units that depend on these applications. The overall protection levels of some major applications were lowered significantly as a result of dependencies on small downstream applications such as job schedulers and other interfaces. For users of applications like the Oracle system, there was an assumption that the environment was adequately protected because the core elements were covered. With greater understanding of actual protection levels, the conversation with the business units went from a very emotional "must have everything" tone to a very objective discussion with empirical evidence that clearly outlined

“From an IT perspective, the business had all the information to make an educated decision regarding any required investments. IT played a trusted advisor role, facilitating the discussion and helping execute the decisions.”

— Dan Inbar, EMC Senior Director of Global IT Technology

“With all the historical resistance surrounding archive and transaction databases, it was significant that we called the process ‘relocation’ rather than ‘archive’. Archive implies that the data is no longer accessible. In our case, archived data needed to be accessible through the native Oracle UI regardless of where it was located. By proving to users that this was indeed the case, EMC was able to overcome the resistance to archive.”

— Bill McHugh, EMC Director of Application Development

the cost/benefit tradeoffs in delivering a high level of protection. As a result, EMC implemented affordable business continuity solutions across those systems that required additional protection and significantly lowered the overall risks the organization faced from potential disruptions.

“From an IT perspective, the business had all the information to make an educated decision regarding any required investments. IT played a trusted advisor role, facilitating the discussion and helping execute the decisions,” added Dan Inbar, Senior Director of Global IT Technology.

Phase II: Applying ILM to Key Application Environments

With a newly consolidated and tiered infrastructure in place, EMC began more detailed examinations of ILM components for two critical application environments: EMC’s Oracle E-Business Suite and its Microsoft Exchange-driven e-mail system.

ILM improves performance and lowers operating costs for key Oracle applications

EMC’s Oracle applications process as many as 1.4 million transactions each day and each transaction can generate more than 3,000 records. With 13-14 copies of production data needed for various processes, from test and development to data warehouse and financial reporting, the growth of transaction data under management was exponential. Despite support staff spending most of their time tuning the application environment, performance problems were a daily issue. For example, users of the sales quoting tool, Direct Express, regularly experienced a seven minute login wait. Applying ILM concepts to the Oracle system, such as ensuring that only the most active transaction data was stored in the production database, would result in a much better user experience and significantly reduce support costs in the process.

Making sure users understand archived information is still accessible

One year after the Oracle E-Business Suite was implemented, the performance issues began cropping up. Root-cause analysis showed the problems occurred as key tables for the production databases grew to 20 million rows. It was clear that an archive process was needed to reduce volumes in the production database, but how to go about that was the key question.

There is often a reluctance to archive because it represents a significant change to the production data set that could potentially affect the integrity of upgrade and maintenance processes. For Bill McHugh, EMC’s Director of Application Development, the answer, in part, came from a change in terminology. “With all the historical resistance surrounding archive and transaction databases, it was significant that we called the process ‘relocation’ rather than ‘archive’,” explained McHugh. “Archive implies that the data is no longer accessible. In our case, archived data needed to be accessible through the native Oracle UI regardless of where it was located. By proving to users that this was indeed the case, EMC was able to overcome the resistance to archive.”

EMC’s DatabaseXtender® (DBX) products were used to automate business-driven policies that defined how and when data should be archived or relocated. These policies ranged from simple to complex depending on the type of data, who accessed it, and the rate at which it was accessed. Just seventeen tables, representing only one tenth of one percent of the total number of tables, but over 80 percent of the overall data growth, were the focus of the archive effort. After a thorough analysis using DatabaseXtender, two-thirds of this data, which included all records six months or older, were relocated from Symmetrix to CLARiON storage.

Users much happier, business more efficient after implementing automated archive

Since using DBX, EMC has kept the growth of its main production Oracle transaction database flat. Prior to the implementation, the IT organization had to spend significant resources tuning application performance just to keep ahead of the challenge. Now, with archive running automatically via DatabaseXtender, IT does not have to tune the application and users are much happier because they no longer wait seven minutes just to login to the application. Most importantly, the business runs more efficiently.

DatabaseXtender has transformed the definition of archiving for EMC. Archived data is still accessible to end users and data relocation has not affected system availability. With more manageable production databases, several important secondary benefits have been realized as well. Now, full copies of production data can be made using EMC's BCV technology, which has eliminated many of the test and development issues that resulted from working with partial production datasets. Migration and upgrade procedures are much faster and more reliable as a result.

ILM facilitates Exchange migration, helps standardize e-mail policies

The next application to be evaluated was EMC's critical messaging infrastructure, a mostly decentralized network of 205 Microsoft Exchange 5.5 servers spread over 70 locations. "The lack of standards across the Exchange network had created numerous inconsistencies that were frustrating users and causing significant backup, retention, and management issues for IT staff that unnecessarily increased risks and costs," explained Chris Mancini, EMC's Director, Windows OS. Sarbanes-Oxley requirements also demanded a more consistent policy for e-mail records across the company that could not be implemented with Exchange 5.5. To make matters worse, Exchange 5.5 was at end-of-life from a support standpoint.

Archive delivers virtually unlimited mailboxes from user's perspective

A decision was made in 2003 to go directly from Exchange 5.5 to Exchange 2003 and consolidate server locations to just two major data centers in order to address these challenges. The new Exchange infrastructure would allow EMC to increase the average mailbox size and significantly enhance the user experience in the process. E-mail retention policies, which allowed older e-mail records to be archived to fixed content storage and ensured compliance with any Sarbanes-Oxley dictated regulatory requirements, were developed based on a user's department or function. The addition of the archive capability effectively eliminated any requirements to manage mailbox size from the user's perspective and create virtually unlimited mailbox storage. Tier 1 mailboxes, including sales and other mission-critical functions plus all metadata, now reside on high-end Symmetrix DMX™ storage protected by synchronous replicas using EMC's SRDF® software. Tier 2 "business-important" mailboxes, covering the majority of EMC's remaining functions, are stored on lower-cost SAN platforms. EMC's fixed content Centera platform holds all archived e-mail records. In addition, multiple node clusters within each data center provide failover capabilities so recovery of the e-mail system is no longer an issue.

With 24,000 of 36,000 mailboxes already converted, EMC is well on its way to achieving much happier and effective e-mail users and much more productive IT support staff. "The great thing about any ILM exercise is that it forces IT, the business, and the users to get together to make business decisions—not IT decisions—about policy," explained Inbar. "The discussions we had with the business owners amounted to mini-classification exercises that helped build consensus and support these decisions. We learned a lot about the ways Exchange was being used, which extended far beyond mere e-mail, and the compliance we gained with Sarbanes-Oxley was an important secondary benefit."

"The great thing about any ILM exercise is that it forces IT, the business, and the users to get together to make business decisions—not IT decisions—about policy. The discussions we had with the business owners amounted to mini-classification exercises that helped build consensus and support these decisions. We learned a lot about the ways Exchange was being used, which extended far beyond mere e-mail, and the compliance we gained with Sarbanes-Oxley was an important secondary benefit."

— Dan Inbar, EMC Senior Director of Global IT Technology

Phase III: Using ILM to Simplify Information Delivery

Unstructured documents increasingly important for critical business processes

Up to this point, EMC had addressed ILM for its most mission-critical structured transaction data and e-mail records. While employees, customers, and partners relied on these records for critical information, they also depended heavily on information such as support bulletins and product presentations that were stored in countless Word, PowerPoint, Excel, video, and audio documents. Most companies find it a significant challenge to manage these unstructured documents. Users often have difficulty finding information, and when documents are available, the information is often inconsistent, conflicting, or out of date.

In 2002 EMC was facing serious challenges getting these documents into user's hands. EMC's Powerlink™ website was the gateway to this information for customers, employees, and partners, but the system had reached scalability limitations. EMC first implemented Powerlink in 1997 as a partner communication website designed to serve a maximum of 500 users. Over time however, Powerlink had grown to support customers and employees as well and was now serving 85,000 users. Unfortunately, focus group feedback revealed that the 85,000 users were not being served particularly well. Users were confused about how Powerlink was structured. They were unsure how to use it and were not consuming the information that EMC was trying to communicate to them.

After completing an initial analysis, the Online Communications organization that supported Powerlink determined that the primary issues could be traced to two factors. First, the original Powerlink was based on communities that defined the information to which a user had access. All users in a community were granted access to the same information. Over the years, the number of communities had grown from five to more than 120. This created significant flexibility and scalability limitations in the system and substantially reduced usability. For example, a customer in a specific vertical market with multiple EMC products was assigned separate communities for the industry they were in and each EMC product group they had implemented. In simple terms, the user had to navigate these separate community pages to gather the information most relevant to them. Second, even though existing documents were required to have associated metadata tags such as owner, abstract, and content type, document owners used these metadata tags inconsistently. One owner would place a text summary of the document in the abstract metadata tags while another would enter a url. These inconsistencies made search results unreliable, which limited users' ability to find documents.

Because of this, the team realized that the current system could no longer effectively deliver information and certainly would not support the continued growth in users and content. Powerlink needed a significant upgrade. So the team initiated a discussion with senior management to explain the challenges. The executives agreed that resources should be invested to rebuild Powerlink, but they noted that thirty other knowledge management initiatives were under review across EMC and asked that an overall knowledge management roadmap be created.

After a more comprehensive analysis, the team discovered that there was no foundation in place that allowed EMC to address enterprise-wide knowledge management requirements. As a result, rebuilding Powerlink became the first phase of EMC's knowledge management initiative. The team's primary objective was to create an enabling foundation that was extensible enough to handle future knowledge management requirements as well.

The team first needed to define a standard way to classify the thousands of pieces of content on the site. Standard document metadata was a clear key to success for this and document owners needed to be held to that standard in order to ensure consistency and quality. The team made a decision to migrate content that was less than twelve months old.

They then began the effort of reaching out to all document owners in order to gather the metadata attributes that would be required in the new system. Any owner that wished to have documents migrated to the new Powerlink were required to update document metadata to conform to the new standard.

The team then made a decision to use Documentum® as the content management system for the new Powerlink. Documentum's content management capabilities and a portal personalization layer, along with consistent metadata enabled a profile- or role-based, rather than community-based, system. This greatly improved flexibility, scalability, and usability of Powerlink. For example, a North American customer in the HealthCare industry with CLARiiON and Documentum products would now be presented the information most relevant to them in a single experience, instead of having to search through multiple communities.

Automated policies help ensure document currency and accuracy

The new Powerlink has introduced the concept of workflows, which enabled the automated enforcement of information policies—a key concept of ILM. These policies cover areas such as document expiration which define how documents are either archived or deleted at expiration. As a result, the content expiration process is much more effective now. Just prior to expiration, the system generates an automated e-mail that is sent to the document owner. Many of these documents are then archived to near-line or long-term storage, or deleted, depending on the governing policy. Powerlink is no longer bogged down by inactive, out-of-date documents.

“Features like workflow and automated policy enforcement are a critical component of ILM. Powerlink workflows ensure that all documents are in compliance with metadata policies at the time the document is created and that authors follow approved templates for specific document types. Now, information delivery is much more efficient. Users easily locate the documents they need and information is up-to-date and consistent across documents.”

— Len Devanna, EMC Director of Online Operations and Strategy

“Features like workflow and automated policy enforcement are a critical component of ILM,” explained Len Devanna, EMC's Director of Online Operations and Strategy. “Powerlink workflows ensure that all documents are in compliance with metadata policies at the time the document is created and that authors follow approved templates for specific document types. Now, information delivery is much more efficient. Users easily locate the documents they need and information is up-to-date and consistent across documents.”

The success of Powerlink is evident in a number of ways. With twelve to fifteen hundred new users added each week, the team expects the current user count of 140,000 to double by the end of the year. Many new employees have noted how up-to-date Powerlink content is compared to the systems in place at other companies for which they have worked. EMC continues to invest heavily in knowledge management with four major new initiatives expected to go live in August 2006. None of these would be possible without the enabling foundation that was put in place for Powerlink.

“Embedded” ILM Has Changed the Conversation between IT and the Business

While EMC is not done with its ILM efforts, a significant amount has already been achieved. Along the way, events did not always happen as expected. Operational realities such as backup reliability often dictated the priorities when logic suggested a different issue—archive for example—was the first order priority. The classification exercise provided great lessons in how a simplified analysis could quickly deliver a majority of the benefits and help avoid the kind of excruciating detail that that often kills the initiatives at customer organizations.

The newly consolidated, tiered, and operationally centralized storage infrastructure has already delivered several million dollars in savings. Organizational changes have been implemented as a result of the knowledge gained throughout the ILM Showcase process. Operational teams have been reorganized along functional rather than platform lines, which has eliminated redundancies and increased efficiency. Backup procedures, for example, are now managed by a cross-platform backup team whereas previously, the backup responsibilities were duplicated within each platform team.

Procedural improvements have been implemented as well. For instance, business units that request new applications requiring storage must now complete an application alignment interview with IT, similar to the one used in the classification effort, as part of the implementation process.

As a result of these changes, the IT organization is no longer bogged down with the reactive support tasks that previously dragged productivity down. Storage utilization has improved dramatically and capacity planning is far more accurate. Users are generally more satisfied with the performance of information systems and the availability of the digital content needed to fulfill their roles. Today, the IT staff supports a greater number of these users efficiently and is able to focus on more proactive, higher value information initiatives as a result.

A common theme across all phases of the ILM effort was that EMC's IT organization has learned to become more business focused. IT staff is now sensitive to users' needs and more literate on the underlying business processes. The business units are far more comfortable in discussing requirements and service levels because IT is speaking to them in terms they can understand. The result is that IT's discussions with the business units are far more objective and less emotional and lead to technology decisions that far more effectively support business processes. Business stakeholders and application owners no longer request boxes, but rather service levels that are linked to their business requirements.

"The various methodologies used throughout ILM force an understanding of the business that most IT staff—especially those supporting infrastructure—would never develop. To me, this is the essence of the value of ILM. All benefits seem to be a direct result of that deeper connection between the technology and the business processes which it supports."

— Dan Inbar, EMC Senior Director of Global IT Technology

"ILM is a great way for IT to align itself with the business by becoming a partner and trusted advisor and helping to enable the end users," said Inbar. "The various methodologies used throughout ILM force an understanding of the business that most IT staff—especially those supporting infrastructure—would never develop. To me, this is the essence of the value of ILM. All benefits seem to be a direct result of that deeper connection between the technology and the business processes which it supports."

"There is certainly no such thing as perfect ILM," explained Pratt. "But even in the real world that is EMC's infrastructure, ILM has led to an IT organization that is far more efficient, flexible, and business oriented."



EMC Corporation
Hopkinton
Massachusetts
01748-9103
1-508-435-1000
In North America 1-866-464-7381

Copyright © 2006 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, EMC ControlCenter, AlphaStor, ApplicationXtender, Captiva, Catalog Solution, Celerra, CentraStar, CLARAlert, CLARiiON, ClientPak, Connectrix, Co-StandbyServer, Dantz, Direct Matrix Architecture, DiskXtender, DiskXtender 2000, Documentum, EmailXaminer, EmailXtender, EmailXtract, eRoom, FLARE, HighRoad, Navisphere, OpenScale, PowerPath, Rainfinity, RepliStor, ResourcePak, Retrospect, Smarts, SnapShotServer, SnapView/IP, SRDF, Symmetrix, TimeFinder, VisualSAN, VSAM-Assist, WebXtender, where information lives, Xtender, and Xtender Solutions are registered trademarks and EMC Developers Program, EMC OnCourse, EMC Proven, EMC Snap, EMC Storage Administrator, Acartus, Access Logix, ArchiveXtender, Authentic Problems, Automated Resource Manager, AutoStart, AutoSwap, AVALONidm, C-Clip, Celerra Replicator, Centera, CLARevent, Codebook Correlation Technology, Common Information Model, CopyCross, CopyPoint, DatabaseXtender, Direct Matrix, EDM, E-Lab, Enginuity, FarPoint, Global File Virtualization, Graphic Visualization, InfoMover, Invista, MirrorView, NetWin, NetWorker, OnAlert, Powerlink, PowerSnap, RecoverPoint, RepliCare, SafeLine, SAN Advisor, SAN Copy, SAN Manager, SDMS, SnapImage, SnapSure, SnapView, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix DMX, UltraPoint, UltraScale, Viewlets, and VisualSRM are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners.

EMC Perspective
H2148