



WHITE PAPER

February 2006

Information Lifecycle Management

Benchmark Study: Maturity Progression

ABSTRACT

Information lifecycle management (ILM) is a sustainable storage strategy that balances the cost of storing and managing information with its business value.

This report is the result of a benchmark study assessing the maturity of over 900 IT organizations in North America and Western Europe. This paper is the second of a two-part series describing the results of the benchmark study; it is aimed at understanding the progression of activities that organizations undertake as they become more mature.

Executive Summary	2
Benchmark Study Introduction	3
Benchmark Overview	3
Overall maturity ..	3
Sample demographics	3
The Maturation Process — Key Findings	4
Balanced performance wins	4
Key elements of information lifecycle management maturity	4
Importance of storage resource management	5
Maturity Model Hypothesis	6
The Maturation Process	7
Overview	8
Differences among maturity clusters	8
Chaotic group highlights	8
Chaotic to reactive	8
Reactive to proactive	8
Proactive to advanced	9
Implications	10
Summary	11
Research Study Demographics	12

Executive Summary

Information lifecycle management (ILM) is a sustainable storage strategy that balances the cost of storing and managing information with its business value. ILM provides a practical methodology for aligning storage costs with business priorities.

In August through October 2005, we conducted a research study to benchmark the ILM maturity of a large sample (900+) of respondents based on our maturity model. The research tested 40 criteria to diagnose each organization's stage of ILM development. In general, the organizations averaged close to the midpoint of the model. Forty-eight percent had achieved a proactive state of maturity, with 34 percent below proactive and 18 percent above. There are significant differences among the maturity levels described by participants from different industries, from different size data centers, and based on the participant's role in IT.

- In general, participants from smaller data centers rate themselves lower than their large data center counterparts.
- Health care and financial services occupy the upper echelon of maturity with average scores significantly higher than the overall average. Manufacturing, services, and information technology are not significantly different from the overall mean. Communications, transportation, utilities, distribution, and public sector are significantly below the overall mean. Clear differences exist, perhaps based on business models and the importance of IT to the business.
- Technical staff participants tend to be more critical of their organizations compared to executives, scoring their organizations below proactive almost twice as often as senior managers.

The bulk of this paper addresses a particular context of the benchmark results, specifically how organizations progress through stages of ILM maturity, and what factors are pivotal in the process. While there are several "lessons learned" from the benchmark study, there are two key lessons in the context of the maturation process:

- It is clear that a balanced approach to achieving ILM maturity is very important. The elements of ILM are highly interdependent; achieving a high level of competence in one without balanced performance in the rest does not lead to success. The key finding here is that the most mature organizations exhibit this balance. Thus the goal should be to achieve a state of balance and mature incrementally from there.
- Of course, it is easy to say "fix everything concurrently," but difficult to accomplish. There are clear "next best actions" at most points in the process of maturation, and some elements are clearly foundational. The benchmark points strongly to granular and resilient infrastructure as a foundation of ILM maturity. It's pretty clear that it is wasteful to invest in alignment, process, and tools without the base capability in IT infrastructure to execute ILM. The same goes for automation. Advanced ILM is dependent on automating storage resource management and removing the need for human resources in data movement and placement activities.

Attaining an advanced state of ILM may result in significant business benefits. The most significant benefits are improved organizational adaptability and improved service to customers, suppliers, and internal end users. The combined benefits of fully mature ILM result in business performance enhancements that include increasing revenue opportunities, reducing costs, and driving competitive advantage.

Benchmark Study Introduction

In August through October 2005, a research project was conducted by the Sun Microsystems Market and Customer Research team to develop a benchmark of performance for the Sun Microsystems Information Lifecycle Management Maturity Model. A total of 905 IT executives, middle managers, and technical staff participated, representing a broad distribution of industries.


The ILM Maturity Model survey is an outgrowth of Sun Microsystems’s work in developing a maturity model to describe levels of growth associated with an ILM implementation. Forty storage management behaviors are tested to complete the survey.¹

Benchmark Overview

Overall maturity

On average, the participants scored their organizations below the center point (45.9 on a 1–100 scale) of the maturity scale, with very few respondents at the extremes. Overall, 34 percent of the participants rated their ILM maturity as rudimentary, 48 percent as moderate, and 18 percent as advanced.

Figure 1. Overall Assessment Rating.



Sample demographics

- The most prevalent industry was information technology (17 percent of participants), closely followed by financial services (16 percent), manufacturing (13 percent), and health care (11 percent).
- Forty-three percent of the participants were senior managers — predominately CIOs, CTOs, and VP levels within IT. Thirty-seven percent of the participants were middle managers in IT and 20 percent were IT technical staff.
- Of the participants, 36.6 percent of the participants had fewer than 100 terabytes of data center disk installed, and only 4 percent had fewer than 25 terabytes; 22.8 percent of the participants had 100–199 terabytes installed and 40.6 percent had over 200 terabytes (22 percent exceeded 500 terabytes).
- About 40 percent of the participants were from Western Europe (evenly split between the U.K., France, and Germany) and 60 percent of the participants were from North America, with almost all of the North American participants from the USA (95 percent).

¹ See page 12 for additional information on the demographics of the participants. For a description of the maturity model and a list of the 40 criteria in the model see the Appendices included in “Information Lifecycle Management Maturity Benchmark Study — Part I — Overview.”

Balanced performance across all information management capabilities is key to reaching the higher maturity levels.

The Maturation Process — Key Findings

Balanced performance wins

There are clear differences in performance, and specific strategies for each group to achieve ILM maturity. Organizations with the lowest levels of ILM maturity appear to be focused on alignment. In the next stage, organizations turn their attention to infrastructure, and continue their emphasis on business alignment and business value integration. In the third stage of development, organizations build their storage management integration and placement engines, providing the linkage between business intent and storage actions. At this point, organizations achieve a large payoff from integrated ILM; costs and risk are contained and information asset management emerges.

Growth beyond that point appears to occur in a balanced manner; all capabilities are developed concurrently in small steps. It is clear that the elements of ILM are highly interrelated; it is almost impossible to accomplish some of ILM's objectives without a broad foundation of multiple processes and disciplines.

High performance groups do everything well, leveraging the integration of business and IT, linking metadata, policy, and classification with automated tools for management and data movement, and deploying on highly resilient, adaptive, and granular infrastructure. These organizations have achieved an optimized state of ILM and are delivering consistently on service level requirements and business value.

Thus, the core message is clear: Seek a foundation level of balanced performance, then grow incrementally toward the ILM vision.²

Key elements of information lifecycle management maturity

While balanced performance is the key to highly effective ILM, several factors appear critical. The factors that are most strongly correlated to overall ILM performance include:

- Granular and resilient infrastructure — Infrastructure, as it turns out, is a key factor in achieving ILM maturity. Only 3.3 percent of participants achieved a proactive state of ILM maturity with less than an average level of infrastructure granularity, and only 6.8 percent of participants achieved a proactive state of ILM maturity with less than an average level of infrastructure resiliency. No participants with less than an average level of granularity and resiliency achieved an advanced state of overall ILM maturity. As a reference point, the same metric for some of the factors with lower correlations to overall maturity was in the 15 to 19 percent range.

Of course it makes sense that infrastructure is a key factor in achieving ILM maturity. Alignment and storage management excellence are somewhat wasted if the infrastructure can't perform the desired actions.

Advanced ILM maturity depends on infrastructure that is flexible as well as granular and resilient. Thus, achieving a high level of ILM maturity appears dependent on virtualization approaches.

- Automated data movement and storage optimization — Those who have attempted to manually manage data movement across tiers of storage or to different infrastructures will clearly understand this need. Due to rapidly changing business conditions and service requirements, optimizing a storage environment without the benefit of automated and policy-based data movement is an incredibly time-consuming and close-to-impossible task.

² See "ILM Vision" white paper at www.sun.com

Both automated data movement and storage optimization demonstrate nearly the same impact on ILM maturity as infrastructure; only 7 percent of participants achieved a proactive state of ILM maturity with less than an average level of maturity in these areas. Only one participant demonstrated an optimized state of maturity without at least an average level of data movement and storage optimization maturity.

- Metadata management — The creation and management of comprehensive and rich metadata is a key enabler of advanced ILM maturity. As with the other factors mentioned in this section, very few participants claimed a high level of performance in overall ILM without at least an average level of metadata management performance.

Metadata management is a key linkage point in enabling storage management automation and automated data movement. Metadata provides the basis for enabling policy by providing information that connects policy statements to specific data objects.

- Content management — Capture, management, and use of fixed content information (archive) is a key ILM service. Again, as with the other factors mentioned in this section, very few participants claimed a high level of performance in overall ILM without at least an average level of archive content management performance.

Rapid growth rates of unstructured and semi-structured information have made some form of content management a critical investment area for most IT organizations.

Managing that information generally implies the ability to “look into” business data objects to determine the business value of the content.

Importance of storage resource management

One of the clear inflection points in ILM maturity is the deployment of a resource management process and tools. After that accomplishment, several things become considerably easier. Optimization, storage resource allocation, linkage of storage resources to business process, and basic measurement and management are enabled.

While some organizations may begin with a “process first” approach using low-sophistication tool sets, it is clear from the study that automated storage resource management (ASRM) yields significant benefits. Accuracy and robustness of infrastructure models, advanced measurement and quality improvement, and automated discovery of changed configurations are all made possible by ASRM.

Storage requirements undergo more rapid change today than ever before. The dynamic nature of storage drives a relatively clear need for ASRM. This is relatively simple; if you can’t measure it, you probably can’t manage it very effectively, particularly if it’s dynamic.

The need for automated storage resource management is clear, as storage and information management requirements change at an increasing rate.

Maturity Model Hypothesis

The ILM Maturity Model consists of five maturity states and five major sections of disciplines.³

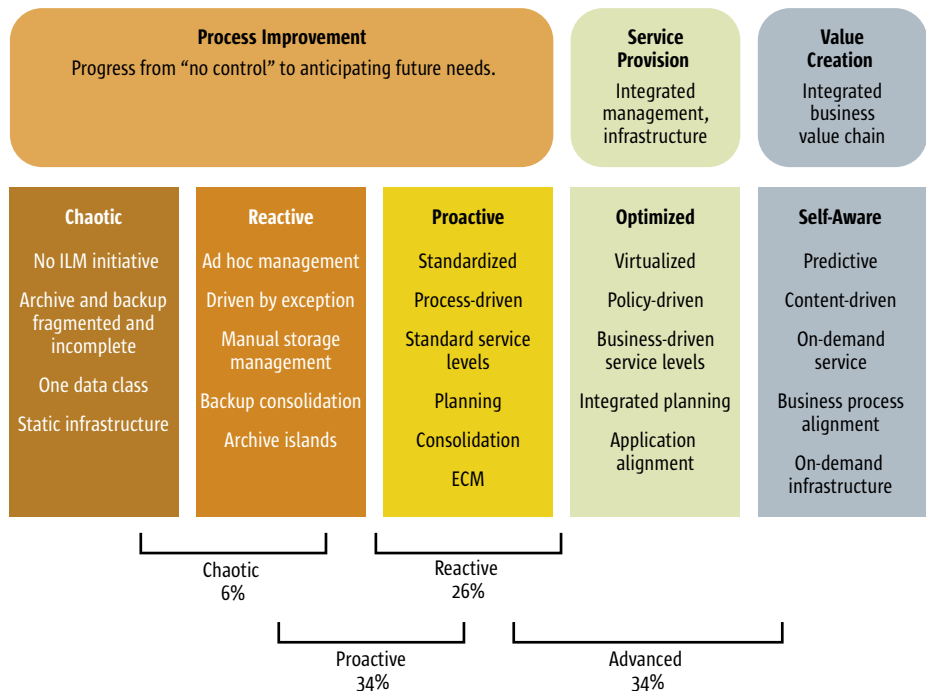
The early maturity stages (chaotic through proactive) are focused on process improvement and have a decidedly IT orientation. The later maturity stages are characterized by a service focus (optimized) and finally a business value creation focus (self-aware).

When we developed the assessment questionnaire, our design point was to define a set of questions that would result in a nice, neat, and normal distribution of results across the maturity model. We expected about 8 percent of the population to fall in a chaotic state, 22 percent in a reactive state, 40 percent proactive, 22 percent optimized, and 8 percent self-aware.

What we found in the benchmark was that models, while instructive, are often “neater” than reality.

In the following section of this white paper we will describe the behaviors reported by four groups of participants that have similar patterns of responses. In general, they map to the maturity model as depicted in Figure 2. The group we will denote as *chaotic* (representing 6 percent of respondents) is composed of predominately chaotic maturity organizations. The group we will denote as *reactive*, representing 26 percent of the participants, is also predominately reactive. A large group (34 percent of participants) with largely proactive behavior will be denoted as *proactive*. And finally a significant group (34 percent) with some proactive and (a few) self-aware participants, but with predominately optimized maturity participants, will be denoted as *advanced*.

Figure 2. ILM Maturity Model, with findings.



³See “ILM Maturity Model” white paper at www.sun.com

The Maturation Process

Overview

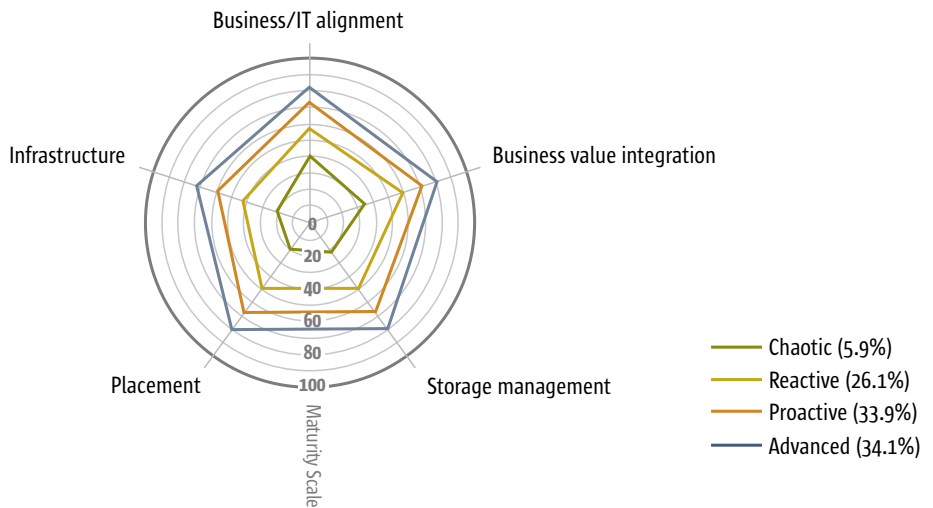
As mentioned in the Key Findings section of this white paper, some interesting observations can be made when ILM assessment performance is analyzed using clustering techniques.

Four key behavioral groups emerge.

The lowest maturity group, which we have denoted as **chaotic** in the chart below, constitutes about 6 percent of the study participants and has an average score that is less than half of the overall mean. The chaotic maturity group is characterized by a significantly higher maturity in business/IT alignment compared to the other elements of the ILM Maturity Model. It appears that the chaotic group has started with a focus on developing relations with their business counterparts and translating that into storage management actions at least at a rudimentary level.

The chaotic group has a proportionally higher percentage of utilities, communications, transportation, and public sector participants and technical staff respondents than the other performance groups.

Figure 3. Information Lifecycle Management Performance Clusters.



The next most mature group, labeled as **reactive** in the graphic in Figure 3, constitutes about 26 percent of the study participants and has an average score that is about 90 percent of the overall average for all participants. Reactive organizations seem to have focused on the same issues as the chaotic group, but with more success.

The reactive group has a proportionally higher percentage of manufacturing, and public sector participants when compared to the other groups. Like the chaotic group, technical staff participants are also more common than in other groups.

The **proactive** group (34 percent of respondents) demonstrates balanced performance, with a moderate focus on alignment and business value integration, but with no section of their maturity more than plus or minus 10 percent from the overall mean score. Overall, the proactive group scores very close to the overall mean. The proactive group has heavy representation from the services and information technology industries. Data centers in the proactive group are also more likely than the overall sample to have fewer than 200 terabytes of disk.

Highlights

Chaotic group

- Focus on business alignment, at a rudimentary level
- Fundamental data protection

Reactive group

- Focus on more sophisticated business alignment and business value integration
- Automated storage management
- Consolidated data protection
- Inflexible infrastructure

Proactive group

- Close to balanced performance
- Consolidated archive and retention management
- Storage management integration focus

Advanced group

- Balanced performance
- Service orientation
- Information asset management competence

Participants in the **advanced** group (34 percent) are a true demonstration of the fact that becoming advanced means “doing most things well.” Due to the interrelated nature of ILM capabilities, it appears to be difficult to become mature without addressing a broad spectrum of people, process, and tools issues. The advanced group has an average score that is 20 percent higher than the overall average, with all but one section of the maturity model within plus or minus 3 percent of the overall mean, and infrastructure lagging at -8 percent. Optimized organizations most frequently come from health care and financial services, and have large data centers.

Differences among maturity clusters

Chaotic group highlights

Generally the ratings on specific criteria for the chaotic group are well below the mean. The group’s highest scores are in the alignment and business value sections of the assessment. Chaotic group highlights include some development of data classification, communications with business process owners, tiered storage, and fundamental data protection. Chaotic organizations appear significantly inhibited by their infrastructure, storage management capabilities, and toolsets.

Chaotic to reactive

The key difference between the chaotic and reactive groups is an increased emphasis on storage management. The reactive group has begun the process of linking business information to storage management. Security, policy management, and measurement process and discipline are emerging. Classification is generally more sophisticated. Data protection has evolved into a consolidated, multiple-level model, linked to service level requirements.

The reactive group has not invested in infrastructure to support its business alignment and integration work; their infrastructure is inflexible. The reactive group scores lower than the mean on all criteria.

Reactive to proactive

The proactive organizations have achieved a moderate degree of ILM sophistication and should be starting to reap the rewards of meeting business information needs at the lowest possible cost. The moderate group exhibits a degree of balanced performance with a slightly less mature infrastructure capability.

Highlighting the performance of the proactive group are advances in metadata management and metadata richness, implementation of storage resource management tools, and sophisticated measurement processes. Measurement includes applications and process metrics, a “single pane” view, and a quality of service/continuous improvement approach. An infrastructure model linking storage resources to user applications is in place.

Information placement has matured significantly. Archive processes have been consolidated and automated retention management has been implemented. There is some evidence that a content management system is driving the archive processes. Balanced advances (compared to the reactive group) have been made in data protection, including automation of recovery operations and further consolidation of data protection services.

The moderates have implemented infrastructure that is fairly adaptive, solidly resilient, and highly granular. The infrastructure is capable of supporting the proactive organization’s advances in storage management and information placement and movement.

Proactive to advanced

Advanced is a proper name for this group. Ratings for all criteria exceed the mean — some by as much as 40 percent and all by at least 10 percent. A “superstars” portion of this group (less than 5 percent of the sample) does everything very well, and consistently exceeds the mean by 2+ standard deviations.

Interestingly, the advanced group, while exhibiting extremely consistent performance across all criteria, remains least sure of its infrastructure. Perhaps this is a result of knowing “how much better” they could be, or it may be that infrastructure has not kept up with the toolset, process, and people, or that the capital expenditure necessary is not yet justified.

Advanced group participants view that they have more work to do in business alignment and business value integration. Notably, they perceive they can leverage their storage management competency to link more effectively to their business information taxonomies, and further broaden the contents of their service level management agreements.

The advanced group has achieved a sophisticated level of ILM. Their policy management, classification, metadata, and infrastructure modeling approaches have reached a point where they are able to enjoy optimized and predictive storage management. They know what’s happening and can proactively make adjustments before storage system events and performance negatively impact users and the business.

They have reached a level of true information asset management and are able to leverage their IT competency to create value for their businesses. They enable the use of historical and archive information to provide improvements in customer service and facilitate advances in decision making. They have substantial confidence in the protection of their information assets, having automated multilevel data protection systems and automated data recovery. Their approach to information lifecycle management consists of a set of services (data protection, storage optimization, virtualization, and archive) that are available to all applications and consistently provide flawless execution.

Base state:

- *Focus on business alignment, at a rudimentary level*
- *Fundamental data protection*

Implications

While the clusters we discussed in the previous section include different businesses at different stages, and do not represent an ongoing study of a set of businesses as they mature, there are some implications that we can synthesize about the maturation process from the cluster analysis.

As noted earlier, chaotic organizations are well below the mean. The group’s highest scores are in the alignment and business value sections of the assessment. Chaotic group highlights include some development of data classification, communications with business process owners, tiered storage, and fundamental data protection. Chaotic organizations appear significantly inhibited by their infrastructure, storage management capabilities, and toolsets. This is the base state on which progress is built.

Important first steps:

- *Rudimentary SRM*
- *Adaptive and granular infrastructure*

In the early stages of maturity, it is clear that IT organizations focus on IT issues and IT infrastructure, and ensure that basic data protection is in place. While maintaining their focus on alignment, the largest strides that organizations make early are to establish some semblance of control, and to make infrastructure improvements. Note that some of these IT operations are quite large (a third have over 200 terabytes of disk), so these issues are not simply tied to small data centers with unsophisticated needs. Something is inhibiting their maturity — either a perceived lack of need, the cost to mature, or the existence of IT management (or platform) silos inhibiting intra-IT cooperation. Some of these organizations may not have sufficient credibility to secure the investment funding required to mature. The primary benefits that such organizations would gain in driving toward ILM maturity would be in risk reduction, cost-effective IT, and IT productivity.

An important step in the early maturity process is the implementation of a storage resource

management toolset and accompanying measurement processes. This may begin as a manual process with unsophisticated measurement tools, but will need to evolve into automated SRM to achieve ILM maturity. It's pretty simple: "if you can't measure it, then you can't manage it."

The importance of granular and adaptive infrastructure is also apparent. It may seem intuitively obvious, but it makes little sense to spend a lot of time trying to understand what the business needs and how to translate that into storage management actions, if you are using 3380 disk and 3490 tape and the devices are directly attached to isolated servers.

Important next steps:

Building the storage management "gearbox" to link business intent to action

- *Metadata and policy management*
- *Classification*
- *Service level management and business alignment*

In the next stage of ILM maturity, a number of foundational disciplines are engaged. Business alignment via service level management, metadata and policy-based management, classification, and initial policy-based management automation are started. These elements are the beginnings of making storage management a "gearbox" that links business intent to management action. The creation of an effective storage management gearbox should drive improved TCO within IT. Storage can be optimized (meaning that storage service requirements and costs are aligned). Storing static data on high-random-performance (and therefore high-cost) disk is eliminated. Inhibitors to taking these next steps include transition costs and most importantly, the initial attempts to secure cooperation with the business. The IT organization is building credibility but needs to take the next step with the business to ensure that it is doing the right things for the right reasons.

As organizations continue to mature, they add more automation, more policy/metadata functionality,⁴ and more sophisticated management tools. Automated SRM (with automated discovery of new components or changed configurations) is added.

Data protection and archive/retention become viewed as ILM "services" and are offered to all applications as such. Data protection schemas increase in complexity, and with policy management, are not a burden on the people managing them. Archive (defined as the capture, management, and use of static, fixed content) is a cross-application service, implying the implementation of a content management strategy. Compliance and governance are integrated into the content management strategy.

Important next steps:

Storage management services

- *Data protection and archive become storage services*
- *Optimized storage*
- *Externalization of storage management*

At this point, the IT organization has an "external" view. It sees the business as its first customer and the customers of the business as potential customers. While IT has not yet evolved into a revenue-creation capability, it is providing solid service to both internal and external users.

The requirement for the business to recognize the potential value of IT is the most likely inhibitor to taking this step. Someone has to take the chance and leverage the organization's IT competency. Once some success is achieved, the ongoing flow of opportunities to contribute to business success may overwhelm IT.

As moderately mature organizations take the next steps to advanced maturity they seem to recognize that they can "do so much more" and appear to be dissatisfied with their service provision role. It's not unusual for sophisticated organizations to be more self-critical than their less-evolved counterparts, and it is interesting to note that some evidence of this attitude was demonstrated by the participants in this study.

Important next steps:

- *Moving to value creation*
- *Process integration with customers, partners, and suppliers*

Advanced organizations, which make up a strong proportion of the participants in our benchmark study (almost a third), and the superstars (one third of that), are on the path from process maturity to service provision to value creation. The ability to secure customer, partner, and supplier relationships to deliver their value inhibits them. They offer their organizations adaptability and information as a business asset.

⁴ See "Metadata Management" white paper at www.sun.com

At an extremely advanced state, information lifecycle management is a fully automated, operational concept. At some point it is like air — taken for granted, until absent.

The primary activities that occur at the final stages of ILM maturity do not greatly involve IT or storage management. IT and storage provide a utility. It is always there. The interesting thing about the superstars in this study is that they do everything about ILM and storage management well. They have moved to the point of ILM maturity where their attention has turned from storage management to information asset management.

Summary

Our purpose in building the ILM Maturity Model was to clarify a roadmap for advancing ILM. We have now taken the next step of testing the model via an assessment and a research study. The assessment will be modified to adjust for the lessons learned in the study, making it a more powerful tool aimed at helping IT organizations decide what to do next on the road to ILM maturity.


Achieving the long-term benefits of a fully mature ILM strategy is dependent on taking the right foundational steps in the context of a long-term vision, and many organizations are making those decisions in the near term. The ILM Maturity Model will help organizations avoid the pitfalls that generally accompany new technology implementations, and will help them to gain solid business benefit from ILM — today.

Perhaps the most important lessons learned from the benchmark study are the importance of infrastructure, data movement/optimization, metadata management, and content management in achieving ILM maturity, and the need to achieve a state of balanced performance. In a nutshell, organizations with an ILM strategy should seek a state of balanced performance and grow incrementally from there.

The largest benefits from the evolution to advanced ILM are business-oriented rather than IT-oriented. Advanced ILM can improve service to business end users, enable the integration of information for reuse in customer- and supplier-facing applications, and improve the adaptability of the organization. Enhanced business and IT alignment closely links business processes and applications to storage management actions and can improve end-user service and the organization's ability to adapt to business changes.

Achieving the benefits of advanced ILM is facilitated by engaging in an orderly maturation process. While a relatively few participants have achieved a value-creation level of ILM maturity, they have demonstrated that it is possible and that balance is the key. They have evolved into an application-aware storage management model that links information management to information use and are deriving significant business value from IT.

Role



Research Study Demographics

The 900+ participants in the sample were distributed across geographies, industry, job role, and data center installed disk terabytes as follows:

Job role


Forty-three percent of the participants were senior managers — predominately CIOs, CTOs, and VP levels within IT. Thirty-seven percent of the participants were middle managers in IT and 20 percent were IT technical staff.

Industry

The most prevalent industry was information technology (17 percent of participants), closely followed by financial services (16 percent).

The 5 percent of the participants grouped as “other” were from agriculture or mining, or did not provide information about their industry.

Industry




Geography

About 60 percent of the participants were from North America (most from the USA). Forty percent were from Western Europe — evenly distributed between the U.K., France, and Germany.


Installed data center disk

Of the participants, 36.6 percent of the participants had fewer than 100 terabytes of data center disk installed, and only 4 percent had fewer than 25 terabytes; 22.8 percent of the participants had 100 to 199 terabytes installed and 40.6 percent had over 200 terabytes (22 percent exceeded 500 terabytes).

Geography



Data Center Disk Capacity



Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN **Web** sun.com



©2006 Sun Microsystems, Inc. All right reserved. Sun, Sun Microsystems, the Sun logo, StorageTek, and the StorageTek logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

JT 0018 A 02/06