

Optimizing the Storage Environment

A smarter approach to information lifecycle management from creation through disposal



Contents

In barely more than a decade, the way we connect and collaborate has changed more than at any time in history. Wireless and web-enabled devices have increased our possibilities for sending, receiving and processing information far beyond expectations. The massive volume of information from these devices—collectively known as “big data”—combined with the need for real-time access is creating an escalating need for data storage. At the same time, these volumes are so large that they are difficult to process with traditional database and software techniques.

The demand for data storage capacity—growing at an annual rate estimated to be as high as 30 percent—is a major driver of escalating IT expense. Costs include not just additional or upgraded devices, but also data center floor space, electricity, HVAC, and ongoing system management.

Smarter systems—instrumented, interconnected, intelligent—are creating a data information explosion. The digital universe is projected to grow from 1.8 zettabytes in 2011 to 72 zettabytes by 2015.

Most organizations, due to over-provisioning and poor visibility into their storage environment, utilize no more than 30 to 40 percent of their available storage. However—although existing storage infrastructure is not fully utilized and not optimized for cost and performance—addition of storage capacity continues to grow exponentially. There are several factors contributing to this paradox. The perception among users is that “storage is cheap,” so storage requests are routinely honored without challenge. Individuals who have the responsibility for performance or access to data are sensitive to performance issues and reluctant to risk degradation of service for want of high-tier storage. Finally, a common contributor to the usage paradox is that, because the provisioning process for storage tends to be labor intensive and time consuming, there is a tendency for application owners and other consumers of storage to over-request. The thinking seems to be that they can save lead time and paperwork by requesting more storage less frequently and letting it sit until needed.



Often, employees who are in a position to approve storage requests also bear the responsibility for performance of the storage environment. They are typically conditioned to approve the highest tier storage in order to ensure maximum performance. A smarter approach to storage management bases each storage request on realistic need rather than conditioned behavior.

A smarter approach to information lifecycle management

Information lifecycle management (ILM) is a process for managing information from conception until disposal, in a manner that optimizes storage cost and access. ILM comprises the policies, processes, practices, and tools used to align the business value of information with the most appropriate and cost-effective IT infrastructure.

The goal of ILM is to place the right data on the right storage at the right time.

Recently, the deluge of data and the need to optimize enterprise storage environments is driving the development of advanced ILM strategies with new processes and services to help maximize the benefits of ILM and minimize the costs of growing storage infrastructures. By combining technology advances (which allow unhindered movement among service tiers) with a streamlined, well defined and automated process flow, IBM's smarter approach to ILM helps ensure that the best decisions are made concerning which resources to use, at the correct points in the process, with the most appropriate people doing the work.

In addressing the storage environment holistically, IBM's approach differs from many vendors in the marketplace. Both storage supply and demand management are taken into consideration. ILM solutions that are restricted to technology alone do little if anything to drive more culturally responsible behavior, with more long-term lasting benefit to the enterprise.

Eliminating the barriers to ILM adoption

The concept of information lifecycle management is not new. IBM pioneered early approaches to records management, which evolved through the data center mainframe environment. In the context of storage management, past ILM efforts have sometimes met with resistance because they tended to be too far removed from the actual delivery of services. An ILM implementation would typically begin—and end—with a vendor performing an assessment and delivering a recommendation for rebalancing data across the existing storage infrastructure. Executing on the recommendations often proved problematic: The work to move data around was manual, time consuming and difficult. And, there was no provision for automatic upkeep.

In the past, there was a perception that the transformation to an ILM environment was difficult to achieve and maintain, and not worth the effort and resources required. Fortunately, newer ILM programs have helped to eliminate the causes for this lack of confidence in the ILM concept and the deployment process.

An additional indicator for the successful adoption of the new, evolved ILM is the alignment of the technology with the business challenges. For example, an enhanced focus on reducing costs requires that the storage pyramid be reversed from a heavy reliance on Tier 1 to a smarter distribution across all tiers. The resulting need for a comprehensive tiering strategy is directly addressed in an intelligent ILM environment.

One of the barriers to previous success with ILM resulted from a failure of heterogeneous storage environments to interact with each other. The remedy proposed by many vendors—migration to a single-source platform—was neither desirable nor realistic for most organizations. New ILM solutions have eliminated this issue with tools that tie the storage assets together to work seamlessly in delivering and managing the environment.

Because adding storage is perceived to be relatively affordable as part of an overall IT budget, the tendency of sellers has been to propose Tier 1 solutions even when it is not specifically necessary. In this scenario, migration to lower tiers (rebalancing) is sometimes difficult. Use of a storage service catalog helps reduce or eliminate over-provisioning at the Tier 1 level by using pre-set standards that permit the use of less expensive tiers during the request stage.

Today, IBM offers a patent-pending storage service catalog used as part of an ILM program, along with the latest tools—proven in actual implementations—which present a comprehensive technical solution that overcomes previous roadblocks to the use or success of ILM in an organization. Without question, one of these previous roadblocks was the very real human resistance to storage service delivery.

Making ILM an operational reality

There are three essential elements to making ILM a reality within a storage environment:

Virtualization improves the utilization and efficiency of the storage hardware resource. With virtualization, storage assets that previously had no common interface can now be used

interchangeably. This means that data can now be moved to less expensive tiers of heterogeneous storage as it ages without interruption to the business. IBM's storage virtualization solution enables seamless replacement of more than 120 multi-vendor storage systems. Optimizing the environment through virtualization generally results in fewer storage components that need to be managed, making it easier to monitor and protect critical data. A virtualization effort in the storage environment can decrease complexity and free up resources, as well as reduce costs.

Standardization, implemented through a storage service catalog, ensures a set of standards that reduces manual intervention and decision making (which takes too long and often results in data being placed in a higher tier of storage than needed). Standardization defines policies for correct size, initial class of service, and management over time. Storage, backup, migration and archival policy definitions can be made easily and incrementally as a company matures. This changes the way storage is requested—from a reactive back-and-forth with the business each day to a proactive, well defined process that is known and understood by all parties.

Automation, with respect to ILM, is analytics driven and can occur in several ways: Storage request workflow automation, workload balanced automated storage provisioning, analytics and policy-driven tier movement within a storage array and across the data center. IBM has focused on all four aspects of automation.

Pre-determined policies for the management of data are key to the ILM concept and can increase storage utilization up to 50 percent. During an IBM Intelligent Storage Service Catalog (ISSC) engagement, data is categorized into specific types and dropped into the correct "bucket." Policies and other criteria size and place the data at the correct tier. Later in the lifecycle of the data, it can be archived, saved to tape or deleted from storage entirely.

The bottom line: Labor and infrastructure savings

The implementation costs for a storage project vary widely depending on the amount of automation involved at the coordination and execution layers. In IBM's own experience with both internal projects and client deployments, standardization alone can reduce the effort by 50 percent. Automation can reduce execution hours by as much as 90 percent.

In a manual scenario, the team meets several times to capture all of the requirements. In an automated scenario, by contrast, the solutions architect furnishes a template to the requester to capture the standardized information included in the storage service catalog.

In addition to the savings in planning time is the savings in time to get a storage request up and running. This reduced implementation time can help reduce or eliminate over-provisioning.

In a recently conducted two-stage pilot for an external customer, IBM used research-developed tools to automatically rebalance 5 terabytes of data based on administrator policies. The result was that a two-to-three-day process was reduced to two to three hours. In stage two, IBM automatically moved 57 terabytes of data in an overnight scenario with no failures.

The right-tiering initiative currently produces \$21,000 per-month savings by using lower-tiered storage instead of higher tiers. Based on 600 terabytes of data for this customer, the savings could be extended to \$2.6 million a year.

IBM can help

IBM has a long history in information management. Today, IBM continues to be a leader in information lifecycle management, offering comprehensive solutions that drive business results and encompass complementary hardware, software and services. IBM storage systems, with enterprise-class disk and tape storage tiers, offer best-in-class virtualization and drive increased return on investment. IBM Research is developing significant new ILM tools that will provide IBM an end-to-end approach—hardware, software and services solutions—unequaled in the marketplace. Additionally, IBM has patent-pending ILM accelerators to provide proven, repeatable techniques for optimizing storage and information management environments. Finally, IBM software has the capacity to enable end-to-end management featuring advanced virtualization, orchestration, automation and robust information management capabilities.

For more information

To learn more about how IBM can help you derive maximum business value through storage optimization and information lifecycle management, please contact your IBM marketing representative or IBM Business Partner, or visit the following website: ibm.com/solutions/itsolutions



© Copyright IBM Corporation 2011

IBM Global Services
New Orchard Road
Armonk, NY 10589
U.S.A.

Produced in the United States of America
June 2011
All Rights Reserved

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corporation in the United States, other countries or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at ibm.com/legal/copytrade.shtml

Other company, product or service names may be trademarks or service marks of others.

References in this publication to IBM products and services do not imply that IBM intends to make them available in all countries in which IBM operates.



Please Recycle
