



Storage Service Management

The Foundation for Information Lifecycle Management

October 30, 2006

Author:

Robert Rogers, Chief Technology Officer, Application Matrix, LLC

Table of Contents

Table of Contents.....	3
Background.....	4
Storage Service Management	4
Features, Capabilities, and Attributes of an ILM Solution	6
Service Level Management.....	6
Business/Service Continuity	7
Availability Management (Operational Recovery).....	8
Availability Management (Data Retention, Disposition, Preservation, and Authentication Policies).....	8
Performance Management	9
Capacity Management	9
Financial Management for IT	9
Incident Management.....	10
Problem Management	10
Change Management	10
Release Management	11
Configuration Management	11
Critical Solution Areas.....	11
Conclusion	13

BACKGROUND

The cost of retaining and managing information can easily become overwhelming; and yet, few enterprises today offer differentiated storage services to their business processes. Understanding the business process lifecycle can redirect resources that are being used by less deserving applications to the “bread-winner” business processes. However, understanding the storage and information aspects of business processes is generally outside the current repertoire of Storage Administrators and Enterprise Architects. This White Paper is intended to discuss the issues and benefits arising from an information lifecycle management (ILM) effort.

Information lifecycle management is much more than tiered storage or even hierarchical storage management. ILM is a dramatic change in information management, focusing on the value of information and delivering storage services based on management requirements. If this sounds reminiscent of the IT Infrastructure Library (ITIL) approach to systems management, or the IT Governance Institute’s Control Objectives for Information and Related Technology (CobiT), they are inextricably linked to one another. Some valuation of the information being managed is a necessity; otherwise, there is the likelihood that everything will be considered “mission-critical.” Equally important are service level objectives that specify how information should be managed and how much information is expected to require management services.

STORAGE SERVICE MANAGEMENT

Storage assets support business processes; it is the business process that dictates availability, business continuity, performance, security, and all other aspects of how storage is budgeted, provisioned, and used. The principles of information lifecycle management help to define those requirements in terms of the ITIL philosophy:

- Service Level Management
- Business/Service Continuity
- Availability Management
- Performance Management
- Capacity Management
- Financial Management for IT Services
- Incident Management
- Problem Management
- Change Management
- Release Management
- Configuration Management

For example, the availability management process focuses on classifying the criticality of business processes, recovery time objectives (RTO), and recovery point objectives (RPO). This information feeds into the architectural design for enterprise storage; prompting technologies

such as local mirroring, hardware or software based redundancy (i.e., including “redundant arrays of independent disks” or “RAID”), remote synchronous or asynchronous mirroring and/or replication, and other technologies. The cost differential between “vanilla” disk storage and highly available, geographically dispersed storage is likely to be enormous. Thus, a classification scheme and service level objectives not only help architecturally, but also financially by helping to build an objective set of organizational priorities that meet budgetary requirements, service requirements, and technological requirements.

Techniques for assessing business processes from the storage perspective are being developed by industry leaders and in several different organizations including SNIA, GGF, DAMA, AIIM, and others. Most of these best practices have their origin in the IT Infrastructure Library (ISO27000), and the IT Governance Institute’s Control Objectives for Information and Related Technology (CobIT).

As previously stated, the features, capabilities, and attributes that would be appropriate for a large, heterogeneous, geographically dispersed, enterprise storage architecture depend greatly on the applications and their criticality. Most commercial enterprises have a wide ranging mixture of business processes, and supporting applications. The most critical business processes usually number a few dozen out of the thousands of applications in a large business. Thus, designing an architecture to support every business process at the same level of criticality either cripples the most critical applications, or inflates the cost of the configuration by an order of magnitude or more by over-provisioning less critical workloads. Information classification is a necessity.

Once information and business processes have been categorized, it is possible to begin applying resources to service those information assets. For the purpose of this white paper we have constructed three categories of service requirements as examples to illustrate the potential architectures that might be employed to service enterprise storage needs. In most real instances, more than three categories are needed, but the examples demonstrate the concept. Note: the environment has been simplified with details such as media refresh, and data format refresh omitted. For data with 100 year retention requirements these would be required considerations.

1. Time-Critical Business Processes

- a. Essential to the operation of the Enterprise. Application outages (planned and unplanned) are constrained to 15 minutes/year. Application recovery time is 15 minutes or less. The recovery point for time-critical business processes provides consistency to the operating system level. One hundred year local data retention with transaction response time requirements of two seconds or less when data are used within a 30 day period. Transaction response time requirements for less frequently used data (30 days to seven years) may extend up to one minute, and data accessed less often than seven years may have transaction response times of up to 15 minutes. After 40 years data will be copied to an offsite archive in secure format with key escrow in a segregated, non-repudiated data transfer. Upon confirmation of the readability of the archived copy, the local copy and all of its backup images will be erased. Storage for Time-Critical Business Processes is requisitioned on a weekly basis. New Time-Critical Business Processes must undergo regression and conformance testing before propagation to the production environment.

2. Mission-Critical Business Processes

- a. Essential to the operation of the Enterprise. Application outages (planned and unplanned) are constrained to no more than four hours per occurrence and no more than one planned occurrence per quarter. Application recovery time is four hours or less. Local site failures also require full service restoration in four hours or less. Local data retention of 20 years with transaction response time requirements of two seconds or less when data are used within a 30 day period. Transaction response time requirements for less frequently used data (30 days to six months) may extend up to one minute. Infrequently used data (six months to seven years) – five minutes; and data accessed even less often may have transaction response times of up to 15 minutes. After 20 years, data will be copied to offsite archives in secure format with key escrow in a segregated, non-repudiated data transfer. Upon confirmation of receipt from the archive facility, the local copy and all of its backup images will be erased. Storage for Mission-Critical Business processes is requisitioned on a bi-monthly basis.

3. Other Business Processes

- a. Essential to the operation of the Enterprise but deferrable in terms of priority. Applications are restored to operational status based on their last backup version with recovery times that may range from minutes, hours, days, or longer. Local data retention based on application and business process requirements. Transfer to offsite archives based on records management (RIM) directives.

FEATURES, CAPABILITIES, AND ATTRIBUTES OF AN ILM SOLUTION

This scenario features an enterprise storage environment consisting of data center resources, and remote office storage resources. The business processes being serviced in this scenario consist of all three previous identified categories (Time-Critical, Mission-Critical, and Other). Servers using this configuration consist of z/OS, UNIX (e.g., Solaris, AIX, HP-UX, Linux, etc), and Windows Server.

Service Level Management

A set of Service Level Management (SLM) processes are a necessity to balance user demands for service and support against budgetary and technology issues. As in the examples of the three classes of workload, priorities are established, expectations are set, constraints (i.e., what happens if the user demand for service exceeds agreed upon boundaries), and penalties for non-compliance (e.g., if availability or performance standards are not met). There are several parts to the SLM process, setting objectives, negotiating agreements, monitoring compliance, and arbitration when there are compliance issues or difficulties in the negotiation process.

In most cases, a SLM process monitors end-to-end response time and other user-based metrics. To augment the SLM process for storage, those metrics must be supplemented with information derived from storage resource managers (SRM) and other storage-centric tools (e.g. Syslog for UNIX systems). In a relatively small number of advanced enterprises, the entire provisioning process for applications is automated and based on service level management, with users

requisitioning resources, obtaining management and financial approvals, generating facilities management authorization for equipment and services, etc.

The ILM Technical Workgroup of SNIA is heavily involved in the design of storage-centric service level objectives that can be used to automate and manage data assets throughout their lifecycle.

Two key considerations are cost and scalability. Service level management is one of the most important contributors to cost containment and scalability. By ensuring that information is accorded only the level of service it requires, resources are conserved and aligned.

Much the same argument can be used for scalability. By ranking and classifying information, the most important, Time-Critical, or Mission-Critical information can get the resources it was allocated, and less valuable information can then contend for whatever is left.

Business/Service Continuity

Time-Critical Business processes should be supported by multiple levels of robustness, including RAIDx, local clustering, metropolitan remote synchronous mirroring, and geographically remote asynchronous mirroring with geographically remote server failover. Metropolitan remote synchronous mirroring is effective across short distances (i.e., about 100km or 62 miles) for both z/OS and other operating systems. Asynchronous mirroring is needed when greater distances are involved to avoid application degradation. For very important applications where close to zero loss of data is required, hybrid mirrored implementations are used as shown by Figure 1. Business Continuity for High Availability Applications.” Synchronous mirroring is used to get mirror data offsite, and asynchronous mirroring then augments the synchronous copy to achieve the distance required to store data out of the danger zone.

Mission-Critical Business processes are also supported by multiple levels of robustness, including RAIDx, local clustering, and geographically remote asynchronous mirroring.

Finally, for truly time-critical data, it is inadequate to remotely mirror to a single failover target. The opportunity for a failure to occur while maintenance is being performed, or worse, after a “live test” while data are still being re-synchronized is too great. When data are irreplaceable, and applications have to resume in minimal time, a “cascading” set of mirrors is required, so that once implemented, there is never an occasion when the primary and secondary site are significantly out of synchronization.

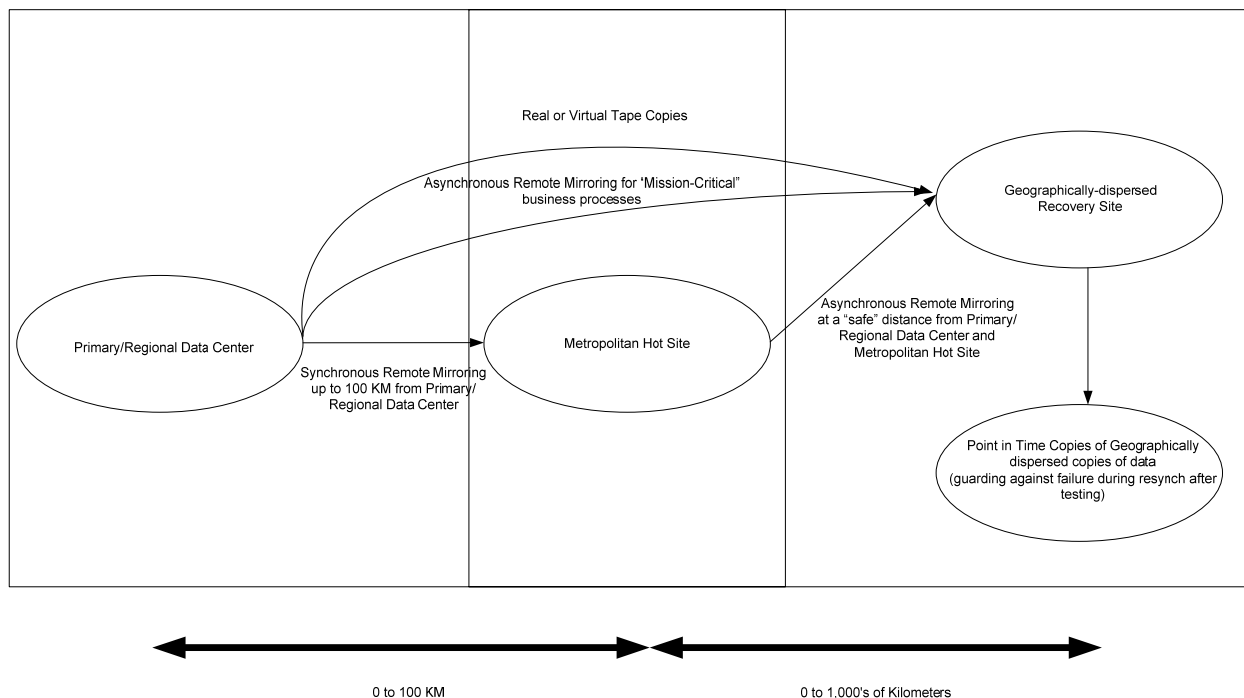


Figure 1. Business Continuity for High Availability Applications

Availability Management (Operational Recovery)

RAID and mirroring solve specific device failure, or site failure issues, but fail to address operational recovery issues. Operational recovery is the condition where data have been corrupted or lost; usually by a processing error (either programmatic or by a user/administrator) and require restoration to a past operational state. Continuous data protection (CDP) solutions provide rollback functions (e.g., by journaling changes at the block, file or object level). For the majority of work, more traditional backup and restore tools do an excellent job of handling open systems backup especially when combined with a virtual tape appliance.

In today’s large IT environments, it is inadequate to create a single backup. At least one is usually required to be maintained off premises. This secondary copy fulfills multiple purposes. First, synchronous or asynchronous mirroring of all the data at a primary site is usually much too expensive (e.g., consider the “Other Business Processes”). Secondly, while mirroring is an effective countermeasure to device or site failure, it provides no protection from a multiple failure scenario in which an application corrupts data which is then mirrored, after which the primary site is destroyed.

Availability Management (Data Retention, Disposition, Preservation, and Authentication Policies)

As previously stated in this white paper, the systems management processes (e.g., ITIL and CobiT) only partially focus on storage. These policies arise from auditability and security requirements and the need for corporate and regulatory governance. In the case of government; they are important from a historical perspective as well. Every business process has these policies; either explicitly defined, or implicit. Data retention and disposition policies are perhaps the easiest to understand; they define the conditions for retaining information, and just as

importantly, how to dispose of the information when it reaches its end-of-life. For example, some legal records must be maintained virtually forever, yet some data age rapidly (e.g., backup versions). In the examples theorized for this white paper Time-Critical Business processes have a 100 year local retention.

Preservation policies usually stem from auditability requirements. For example, auditors might dictate that records related to a business process be retained at the conclusion of a quarter or fiscal year. A simple copy of that information is probably not sufficient; instead, a point-in-time copy, written to non-volatile media, with metadata describing contents, purpose, retention, and authorization may be the only way to satisfy corporate and regulatory requirements. Thus, preserving information in a particular “state” or condition is a very common need for IT. Both Time-Critical and Mission-Critical data will probably migrate into an archival storage system with restricted access, write-once, read-many, characteristics to ensure authenticity and preservation. There are a variety of archival storage products in the market today. Approaches to data preservation differ from product to product with only a few taking into consideration the infrastructure requirements needed to be capable of *productively* using data years from now. The Library of Congress, NARA, as well as other academic groups such as the International Research on Permanent Authentic Records in Electronic Systems (InterPARES) are doing work in this area.

Authorization issues are especially important in storage management because of the need to restrict authority (partially because the consequences of an error may not be reversible) and ensure audit trails for every action. Authorization to use data should not be automatically granted based on the user’s ability to create data. For example, when a process creates reports for auditors, the auditors (and not the creating application) are the owners of that information.

Performance Management

Performance management functions as support to the other systems management disciplines. When responsiveness issues are reported, performance management tools identify the bottleneck. Performance management monitors and tracks service level compliance to correct deficiencies in service or unanticipated peaks in workload. The performance management discipline also supports Financial Management by validating that bottlenecks are being caused by resource shortages that require additional hardware or software and not simply re-distributing workload priorities.

Capacity Management

Most IT environments are constantly growing; and according to industry analysts, storage usage is growing at a much faster rate than most other IT resources. Monitoring and tracking resource utilization so that surprise is minimized and requisitions for additional storage resources are anticipated is the primary role of the storage capacity analyst. Element Managers provide monitoring and historical information for discrete vendor hardware and software. Efforts underway through the Storage Networking Industry Association to advance the Storage Management Initiative (SMI-S) are leading to broader, more vendor agnostic solutions.

Financial Management for IT

The financial management role in IT is very important. To avoid outsourcing (just one of the

financial issues in the corporate sector), it is critical for IT to justify its economic value to the enterprise. Storage has historically done a poor job of demonstrating its value. ISO standards, ITIL, CobiT, and other work (especially in information lifecycle management) are helping IT management to articulate the value it provides to the enterprise. Consider the issue of information availability. In this white paper we have theorized Time-Critical Business Processes with a 15 minute recovery time objective (RTO). How an organization calculates the cost per minute of downtime has a significant effect on the economics of an enterprise storage architecture. The investment necessary to substantiate the need for 15 minutes or less of downtime is very large for most time-critical business processes with primary, secondary, tertiary disk, and large capacity communication circuits.

Financial management is a much broader area than simply chargeback; although that is one key focus area. Financial management includes information valuation, risk management, and other economic considerations. Very little analytic research has been done in the industry to build economic models; research in this area is just beginning.

Incident Management

Incident Management as it pertains to storage affects information security. Monitoring and auditing logs for unauthorized data access, and monitoring element managers for breached security rules. Most of the element managers supporting enterprise storage hardware have their own, independent, authentication and authorization schemes.

Problem Management

Problem determination (PD) and problem source identification (PSI) are needed to appraise issuing affecting service continuity and service quality. Generally, the problem management discipline feeds the metrics that constitute service level management. For example, the Time-Critical Business Process is expected to maintain availability with a service quality of two seconds for all but 15 minutes per year. To ensure those objectives are met requires reporting and identifying the problem's source for every service issue, scheduling a correction, validating its efficacy, and building a regression test for subsequent releases to ensure such a problem does not recur.

Change Management

Change management policies and procedures tend to reduce costs, improve service quality, and also improve information security. The cost reduction aspect is obvious; by assessing the risk of each change, and validating that institutionalized procedures for testing changes are followed, system quality is improved. In the commercial sector, before any change is allowed to affect systems running Time-Critical Business Processes, the changes generally need to go through several increasingly rigorous testing processes. In the most sensitive applications, it may require months of testing changes (microcode, new procedures, new software levels, etc) with many risk assessments and practice "backout" exercises before changes get propagated to highly available environments. In many of these cases, special monitoring procedures are instituted to observe particularly risky changes so that any service degradation is minimized.

Release Management

Release management is important to enterprise storage systems from the perspective of managing the environment; however, when long-term archive is involved (e.g., Time-Critical and Mission-Critical Business Processes), there are special needs. Release information is vital to reconstituting an application once it has been archived for some period of time. The importance of preservation metadata describing the data, program, infrastructure, and hardware co-requisites all need to be assessed with every new release and function. Information may need to be retrieved from the archive, migrated into new release format, checked for authenticity between old and new formats, and re-archived as a new entity beside the original. Few vendors have the metadata necessary to ensure that long term archives actually live up to their service level requirements.

Release and Configuration management disciplines intersect when there is a high level of effort to reduce security vulnerabilities. In such environments, new releases are registered and deployed, and configuration/asset management tools monitor vital product data (e.g., executables, registry entries, and configuration specifications) for unauthorized changes. The Enterprise Storage Administrator will usually participate in the release management process to help ensure compliance with release policies.

Configuration Management

A combination of hardware, facilities, and portfolio management make up the discipline called configuration management. Hardware and facilities management constitute the physical inventory components; portfolio management constitutes the software component.

In the case of portfolio management it is important to be capable of “registering” a baseline or authorized configuration and monitoring its status to ensure that only authorized configuration changes are made.

Configuration management affects most of the other systems management disciplines because without an inventory of resources it becomes impossible to align resources to service requirements.

CRITICAL SOLUTION AREAS

The Aberdeen Group’s July 2006 study entitled “The Information Governance Benchmark Report - *A Needed Strategy for the Enterprise Backed By Viable Solutions*” indicates that the chief concern among 115 enterprises was risk analysis and management (42%) followed in relatively close succession by the need for automated processes (34%), aligning IT policy risks and operations (32%), and documented policies and procedures (29%). These concerns all point to a need for highly developed service objectives, policies, and procedures. To create such processes, requires that business needs and resources be aligned to one another ensuring that information and associated processes are managed efficiently and effectively. In several key areas, the study’s responders re-affirmed that policy-based management (i.e., automation) was exceedingly important to their success at information governance. About two-thirds of the responders indicated that enterprise-wide policy management was “very important” for both structured and unstructured data.

Perhaps of even greater interest was the report’s observation that only 22% of the companies had a set of key performance indicators (KPIs) for information management. The lack of KPIs indicated that these companies have recognized their need and are acting to remedy the situation.

Beyond the need for classification, policy management, automation, and monitoring tools, the business processes theorized to exist (“Time-Critical,” “Business-Critical,” and Other) have a requirement for rapid recovery, offsite storage, and redundancy which imply high speed networking connectivity between geographically separated locations. Table 1 lists the primary components of an enterprise storage architecture starting with management elements and working into the physical elements.

Component	Component Description
Information Classification	Identifies information, its origins, and usage. May use one or more schemes for information classification including: <ul style="list-style-type: none"> • Metadata-based • Content-based • Business Process Mapping
Taxonomy and Ontology	Organizes classified information
Information Valuation	Applies risk and valuation metrics to classified information
Policy Management/Automation	Applies data and information management rules to classified information to align the information with resources for: <ul style="list-style-type: none"> • Quality of Service • Provisioning • Retention • Migration • Archive • Backup • Recovery point synchronization • Synchronous and Asynchronous remote mirroring • Custody Management (non-repudiation)
Storage Resource Management	Monitors and alerts on exception conditions in the storage configuration
Authentication/Authorization	Access to information and authorization to create/delete/update/move/encrypt/decrypt
Capacity Management	Monitors resource utilization and projects resource allocation requirements for budgetary purposes
Problem Management	Links operational issues to storage management for resolution and reporting
Change Management	Links operational changes to storage management for risk assessment, monitoring, and reporting
Release Management	Links operational changes to storage configuration management for monitoring, assessment, and reporting

Table 1. Solution/Component Description

The most important design principle is that storage is a resource that can and should be managed from a service-level perspective. The day-to-day needs of storage service management require a

much higher semantic than spending hours, days, or weeks, creating LUNs, zoning SANs, and presenting LUNs to servers. Twenty years ago, it was commonplace for a storage administrator to be responsible for ten gigabytes of storage. Today, a storage administrator is typically responsible for 20 terabytes or more of storage. Five years from now, each storage administrator will probably be responsible for one half petabyte of storage (500 terabytes) or more. To scale enterprise storage to meet enterprise needs; it must not be a low-level function that requires manual intervention at every step. It must be a service with multiple levels of authorization, service level management, and automation. That “complete” solution is not fully realizable today, but by designing for that eventuality, it is possible to achieve a large measure of the benefits, even with today’s technology.

Virtualization is one of those technologies available today to help to mask environmental changes from end-user applications. Although there is an initial cost, virtualization reduces the amount of change necessary to implement new technology. For example, an enterprise backup tool need not be aware of the underlying physical hardware in an enterprise storage environment if virtualization is employed (at either the front or the back-end of processing).

CONCLUSION

Today’s storage environments have to continue operating; thus, any new design must be incremental; with the ability to apply policies, procedures, and new technology without major upheaval. In highly dynamic environments where change is constant, a new architecture might be introduced and business processes migrated by attrition. However, in most instances, the lifecycle of a business process or software application is much too long to contemplate the attrition methodology. Storage service management is an on-going process because of the constant change within the enterprise. As it exists today, it is mainly a manual, labor-intensive, process, requiring detailed application, data center, and business process knowledge. In the future look for analysis, automation, and policy management tools to simplify the task of storage service management. An enterprise is nothing without its data; storage service management helps to ensure that information is available, timely, and within budget.

About the Author

Bob “Mister” Rogers has more than twenty-five years of storage, systems, and performance management experience. Mr. Rogers is presently Chief Technology Officer and founder of Application Matrix, a startup focusing on Information Lifecycle Management (ILM) and Business Process Management (BSM) tools. At BMC Software, he was Chief Storage Technologist of the Storage Division. As Chief Architect at SOFTWORKS, he orchestrated the development of storage software for heterogeneous platforms and was instrumental in leading the company through its 1998 IPO and EMC's \$192 million dollar acquisition in 2000. Mister Rogers led several of IBM's product introduction programs including ADSM (now the Tivoli Storage Manager), HSM, and DFSMS (IBM's Systems-Managed Storage) prior to joining SOFTWORKS. Mr. Rogers is a prolific writer, speaker, sought-after consultant in his field, and one of the founding members of SNIA's ILM Technical Working Group.

About the ILM Professional Services Task Force

This SNIA DMF ILM Professional Services Task Force (PSTF) is a collaboration of Professional Service organizations (large and small) from across the storage and information management industry with the intent to reduce confusion in the ILM market by providing a common lexicon, producing educational materials for end users, and setting expectations among users about what is possible from ILM Professional Service engagements.

About the SNIA

The Storage Networking Industry Association is a not-for-profit organization made up of more than 300 companies and individuals worldwide spanning virtually the entire storage industry. SNIA members share a common goal: to set the pace of the industry by ensuring that storage networks become efficient, complete and trusted solutions across the IT community. To this end, the SNIA is uniquely committed to delivering standards, education and services that will propel open storage networking solutions into the broader market.