

WHITE PAPER
February 2006

Best Practices in Information Lifecycle Management Security

ABSTRACT

Data security is a critical component of information lifecycle management maturity. This paper describes best practices in securing data from the perspective of information lifecycle management (ILM). It is intended as a high-level introduction to the main categories of storage security and considerations in balancing conflicting storage management priorities.

1	Executive summary	2
2	Introduction	3
3	Major ILM security categories	4
3.1	Physical security	4
3.2	Access control	4
3.3	Encryption	5
4	Implementation and improvement considerations	6
4.1	Overall considerations	6
4.2	Physical security considerations	6
4.3	Access control considerations	7
4.4	Encryption considerations	8
5	Conclusion	9

1 Executive summary

As organizations adopt business models that provide employees, partners, suppliers, and customers with greater access to sensitive information, they create increased business opportunities as well as increased risks. Escalating cyber terrorism, data theft, and evolving corporate governance issues have intensified the need to secure corporate information. Organizations increasingly understand this need and are responding to it. However, a recent survey conducted by Sun Microsystems shows most senior managers and technical staff rate their security processes as not fully integrated with storage management policies¹. Organizations with medium-to-large data centers (at least 100 terabytes of disk) score themselves significantly higher on this integration than organizations with smaller data centers. Organizations with large data centers (at least 200 terabytes of disk) rate themselves higher at linking storage security policies with network security policies. Technical staff, who are more likely than other respondent groups to have a detailed understanding of their organization's current situation, rate storage and network security policy integration as low. This is troubling, given the recent massive efforts to create networked storage environments.

Secure information lifecycle management that protects information dispersed across a diverse environment requires an overall solution that is comprehensive and scalable. No single vendor offering comes close to meeting these needs. Data security markets today are dominated by point solutions, and enterprises should not expect significant out-of-the-box integration for years. However, point solutions for physical security, access control, and encryption, when coupled with appropriate IT and business processes, can offer real benefits now.

Enterprises can start by identifying critical data security requirements and including them in their data classification processes. Data users can be categorized on the basis of the needs associated with their roles. An appropriate security approach can then be designed that incorporates IT processes and vendor security offerings that will support and strengthen them. It is important that organizations select security solutions based on the real risks they address. Other best practices include:

- Using established standards and models as references in meeting the organization's unique security needs
- Focusing on user productivity first to gain strategic advantage through access to needed data
- Protecting data from theft, mutilation, unintended disclosure, or deletion wherever it may exist, whether "at rest" or "in flight"
- Creating multiple layers of security, without creating excessive management complexity
- Ensuring that security processes are incorporated into overall business and IT processes

Each organization will need to construct its own unique storage security solution, which must continue to evolve, adapting to new opportunities, threats, and capabilities.

¹ See "Information Lifecycle Management Maturity Benchmark Study: Overview" and "Information Lifecycle Management Benchmark Study: Maturity Progression" at www.storagetek.com/solutions/white_papers.html.

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.

Storage security focuses on protecting data from unauthorized access, destruction, or manipulation.

As enterprises migrated from direct-attached storage systems to networked storage environments, they exposed data to new security threats.

Copies of data are increasingly spread out across multiple mobile devices and partner sites.

2 Introduction

Information lifecycle management (ILM) is a sustainable storage strategy that balances the cost of storing and managing information with its business value. ILM strategies that are well-executed should result in a more agile organization, reduce business risk, and drive down both storage unit and storage management costs.

Storage security from an ILM perspective focuses on securing data from unauthorized access, destruction, or manipulation. This focus includes preventing unauthorized people from accessing any content, and allowing trusted people to access only the content their role requires and preventing them from misusing or damaging that content (maliciously or otherwise). ILM adoption efforts without adequate storage security can damage data integrity and increase the risk of disclosure to inappropriate parties. Nevertheless, knowledge worker access to data is fundamental to creating strategic advantage. Successful storage security design and implementation must balance the critical (and conflicting) requirements of access and protection.

The discipline of storage security has become increasingly important in recent years. Enterprises allow data access from outside their traditional perimeters, broaden access by implementing networked storage, and create potential exposures by using IP-based storage networks. Changes in government regulations are also intensifying the focus on storage security.

Data needs to be properly preserved, whether it is archived within storage systems or actively in use inside or outside the data center. It is important to focus on security threats both external and internal. It is also essential for organizations to consider the business value of their data in establishing or enhancing their security practices, though many organizations fail to do so.

Storage networks have matured to a point of complexity that requires additional perimeter, as well as internal, security services to ensure data integrity. As enterprises have migrated from direct-attached storage systems to networked storage environments, they have logically centralized their data, but exposed it to new security threats by providing wider access to the storage array.

New threats have also been created as additional copies of data are being downloaded to partner sites and mobile devices. Furthermore, the use of disaster recovery approaches, such as replication, remote disk, and remote tape, has increased the vulnerability of the data at rest in remote sites.

This white paper examines the three major ILM security tactics to address these issues: physical security, access control, and encryption. While these topics are broken out individually, they are closely interrelated. Successful ILM security is dependent on managing these relationships.

3 Major ILM security categories

3.1 Physical security

The security of media, hardware devices, and data centers, plus the data flowing between each of these, all must be considered.

Physical security, from an ILM perspective, relates to the safeguarding of data in all its forms and states. This includes securing data both while it is “in flight” and “at rest.” The security of media, hardware devices, and data centers, plus the data flowing between each of these, all must be considered.

At the media level, loss and corruption are the biggest risks. While loss of removable media (whether tapes, removable disks, or some other format) is a risk, physically transporting data can offer security as well as other advantages over network data transmission. Encryption can be added to any of these techniques, offering additional safety as well as increased cost and complexity. Media with write once, read many (WORM) capabilities provide another means of achieving physical security.

Physical security and access control at remote or disaster recovery sites must be as rigorous as at primary data centers.

Hardware devices also require physical security commensurate with the data they may contain over time. The concept of creating logical “gaps” between storage devices used for different purposes is worthy of consideration, despite the additional costs and productivity hindrances they may produce. Gaps created to separate different levels of use are less likely to create productivity issues. For example, having a gap between a data warehouse (central repository) and more accessible — and thus exposed — data marts (distributed functional copies of data) can improve security.

Physical security is far more effective if actively monitored.

At the data center level, perimeter security is of primary concern. However, with current technologies the concept of a “perimeter” can be vague. Physical security and access control at remote or disaster recovery sites must be as rigorous as at primary data centers. These sites must be periodically inspected for compliance with security policies.

Historically, storage networks have been difficult to secure. Because data cannot be fully secure unless the storage network environment is secure, this has meant avoiding the use of networks for sensitive data or limiting their use to encrypted data. Recently, however, great strides have been made to improve the security of physical networks.

While all of these layers of physical security reduce risk, active monitoring is required to provide fully effective physical security. However, analyzing alerts and investigating incidents is very labor-intensive and is often not done. More intelligent tools and processes are needed to facilitate effective monitoring and improve overall physical security.

3.2 Access control

Access control can be defined as the rules that govern who can see or modify information. There are many different potential users of an organization’s data, including customers, partners and suppliers, regulators, auditors, and prospects. The “castle defense” mentality that assumes anyone within the organization can be trusted and all others cannot has no place in today’s business environment. Role- and policy-based controls enable the building of complex rules that allow various types of access to different classes of data, assuring security and providing affordable productivity.

Access control can be defined as the rules that govern who can see or modify information.

Basic access controls restrict read, write, and delete privileges based on a user’s identity and successful authentication. Access controls restrict basic access to an object and can be implemented in the file system or database system. Third-party products, applications, and some storage and operating systems can provide the manageability advantages of role-based access controls. Note that access controls protect an object only within controlled systems. Once data is downloaded or otherwise moved outside the controlled system, other tools or techniques must be adopted to maintain control.

Logical controls further enhance content security and support access controls, especially in structured data systems and applications. Security can be significantly enhanced through proper database and application design, and effective use of structured database management system (DBMS) features. For example, referential integrity supports the soundness of the data and enhances access controls. All major DBMS products allow administrators to assign restrictive access to tables while granting wider access to table views, an essential security mechanism. DBMS products include additional vendor-specific security features, such as data labeling for multilevel security, triggers, and row-level access controls, but applying these to established databases may impact performance.

3.3 Encryption

Encryption is a process of transforming data into a format that only the intended recipient can understand. Should unauthorized access occur, encryption would prevent the intruder from reading or manipulating the data. Encryption converts data into cipher-text, which can only be accessed through appropriate credentials or keys. This technique is particularly useful in situations where it is impractical to prevent unauthorized access to data — for instance, while it is in transit across untrusted or hostile networks.

There are three commonly accepted layers and approaches to encryption. Each layer approach delivers an encrypted solution in a different manner and addresses different requirements.

- Application layer
- Database or file layer
- Storage layer

Implementing encryption in the application layer is known to be problematic. Application-based measures can require extensive coding changes, create inconsistencies across systems, and produce ongoing maintenance headaches.

Database-layer encryption allows enterprises to secure data as it is written to and read from a database. Database-layer encryption will also secure data in the file system that the database is using to store the database information. Ideally this type of deployment is done at the column level within a database table and, if coupled with database security and access controls, it forms a sound policy to prevent theft of critical data. Database-layer encryption secures structured data against a wide range of threats, including storage media theft, database-layer attacks, and compromised database administrator access.

Storage-layer encryption enables enterprises to encrypt data either at the file layer (in network-attached storage or direct-attached storage) or at the block level in storage area networks. This type of encryption is well suited for encrypting files, directories, storage blocks, and tape media. In today's large storage environments, storage-layer encryption addresses a requirement to secure data without using logical unit number masking or zoning.

Enterprise digital rights management (DRM) uses encryption to provide granular controls tied to the individual object. Unlike basic encryption, enterprise DRM follows the object throughout its life cycle, but requires deep integration with the business infrastructure and applications, and may not work between different organizations. Example controls include read, edit, forward, copy, paste, delete, or expire the file after a set time.

Encryption prevents unauthorized access or modification by converting data into cipher-text.

Database-layer encryption protects against a wide range of threats, including storage media theft, database-layer attacks, and compromised DBA access.

Storage-layer encryption enables encryption at the file layer (NAS/DAS) or at the block level in SANs without LUN masking or zoning.

Enterprise digital rights management provides controls tied to the individual object; it follows the object throughout its life cycle.

Supplementary security has some advantages, but can cause undue management complexity.

Security classifications should be based on risk assessments.

4 Implementation and improvement considerations

4.1 Overall considerations

The old adage that security is no better than “the weakest link in the chain” is not entirely true. An attempt to breach security would usually take time to find (or randomly hit) the weakest security link. So supplementary security, to the degree that it doesn’t create unreasonable management complexity, has some advantages. Too much complexity probably means increased cost and vulnerability as the organization struggles to monitor its own weaknesses. Perpetrators will be looking for these weaknesses as well, and thus it makes sense to balance capability and complexity in adding tools, techniques, and processes to a storage security approach.

Data classification by business value is an integral part of the ILM process. Data classification by security value should occur as part of the same data classification exercise. Security classifications may be externally dictated independent of data’s “business” value due to compliance requirements. Absent externally dictated requirements, security classifications should be based on risk assessments that take into account the likelihood and impact of inadvertently disclosing or losing data, or losing the ability to verify the validity of data. Even externally dictated compliance requirements, which are often not very specific, should be evaluated based on the risks associated with noncompliance.

The design of many of today’s *computing environments* is hindering rather than helping *knowledge workers in their most basic work*. An effective storage security component of ILM should improve the overall business value of data by balancing better knowledge worker access to data with appropriate risk-based security measures.

4.2 Physical security considerations

Organizations should weigh real risk and “weakest link” concepts in evaluating the appropriateness of available enhancements.

There are many actions that can be taken to improve ILM security at each of the physical layers examined. As we highlight these actions, it is important for each organization to evaluate their applicability in the context of balancing capability, real risk, and the “weakest link” concept.

At the media level, efforts should focus on preventing loss of media, unauthorized reading of media, or alteration of media, especially undetected alteration. Loss prevention can be enhanced through better perimeter security. For example, using a locked tape library, versus unsecured drives and vaults, hinders the ability of someone to walk off with or misplace a tape. Encryption is a common supplement to the physical security of media, and helps prevent unauthorized use of media in transit. However, this requires careful processes supporting the long-term protection and retention of encryption keys.

Using WORM media is a common means of preventing data alteration.

Using WORM media is a common means of preventing alteration. While this technique helps to prevent accidental and malicious data deletion and modification, it often needs to be used in conjunction with other processes that track media, if proof of assured authenticity or avoiding legal repudiation is important.

Media and hardware should be destroyed or completely wiped as part of a controlled decommissioning process. Without this step, it may be possible for the next owner or user to retrieve “deleted” data, resulting in security and privacy violations.

Data cannot be fully secure unless the storage network environment is secure.

At the data center level, there are highly effective, simple security measures that should not be overlooked or allowed to lapse into complacency. For example, closed-circuit TV cameras can be highly effective in prevention, solution, and prosecution if carefully used.

Activity detection and monitoring can identify and record unusual or unapproved activity.

Physical access to the data center — and even areas within it — should be tightly controlled. The access list should be reviewed periodically for changes in roles.

Data cannot be fully secure unless the storage network environment is secure. Best practices in storage network operational security dictate that the fabric or network should provide authentication, test authorization, ensure transmission integrity, protect data privacy, and highlight and survive attacks.

Activity detection and monitoring can identify and record unusual or unapproved activity in applications, on servers and workstations, in storage, and in database management systems, often generating immediate alerts. Activity is usually monitored using a network sniffer, an agent or near-real-time log analysis of the host's audit logs. Native database auditing is included in all database management systems, but it can dramatically degrade performance depending on the depth and breadth of auditing applied. Third-party tools resolve some of these performance issues while including advanced alerting features based on user behavior. For example, an alert would be sent to a security administrator if a database administrator runs a query on credit card numbers.

Although it's not a "silver bullet," well-designed and properly used storage management software can go a long way toward ensuring physical security. Unfortunately, the management interface itself can pose a security risk. If an unauthorized user with ill intent were to breach the system, he or she could use the management interface to reassign storage resources, redefine policies, or otherwise compromise the security and integrity of data. Because significant damage could occur if an unauthorized user were to gain access to the management software, it must be used in conjunction with carefully controlled, role-based user IDs and passwords.

Procedures should be centralized and standardized to help meet reasonable user needs.

4.3 Access control considerations

As a general rule, access control *procedures should be centralized and standardized* to help meet reasonable *user* needs. Although centralization and standardization, if not done correctly, can facilitate greater *damage* from *unauthorized* access, poorly managed islands of *storage* are even more likely to create *vulnerabilities*. The use of role-based security controls is critical to ensure that user *access is facilitated, but only to the degree of reasonable need*.

Data users should not be granted access until they have signed a security agreement.

Data users should not be granted access until they have signed a security agreement. This document provides education and an incentive to users to secure data that is downloaded to files on desktop or mobile computers, or to those generating reports, such as HTML documents, that may be widely dispersed.

The ad hoc nature of users downloading data to a local DBMS or file significantly compounds the challenge of securing the data.

We recommend that for all access requests, applications pass the original user ID to the DBMS, providing for better security auditing and usage tracking in the database. This model requires that authentication is performed outside the DBMS (by the network, operating system, or specialized security software), or the database administrator must maintain passwords in the DBMS. External authentication is more desirable because it removes password maintenance from the database administrator and places it with a dedicated security administrator function.

Policies should state whether read access to archives requires auditable records. If required, access methods should record evidence of read access in a secure audit log. Write access to archives should be prohibited except for initial data load operations. Delete access also should be policy-based and recorded in a secure log if policy dictates.

Using internal DBMS security features, a user authorized to access a data warehouse can be limited to accessing specific tables and columns of data within the data warehouse. However, a lot of manual effort is required to translate the security rules and policies of the data warehouse to a subset of the data (and related schema) for a data mart. Many of the business intelligence tools that allow data to be stored for further analysis in a specialized data structure (such as a multidimensional DBMS) provide thin security at best. Often, those DBMS technologies are maintained by users and IT personnel who may not understand or care about the security issues facing an enterprise. Without applying security rules to this data, an authorized user could inadvertently leave it unsecured. The ad hoc nature of users downloading data to a local DBMS or file significantly compounds the challenge of securing the data.

Time- or location-specific restrictions on data requests should be strongly tied to the identity of the requestor. Therefore, the ability to verify the identity of users, devices, and applications through an appropriate authentication method becomes a critical component of securing data within an enterprise.

4.4 Encryption considerations

Although encryption is often essential, it is not a total solution. Most enterprises leave their systems vulnerable to attack regardless of their use of encryption. Depending on the storage function, the level of security required, and the trust zones along the storage path, data encryption can be implemented in balance with other security efforts in several ways.

Encryption keys are used for proper and secure identification among entities involved in data exchange, as well as to encrypt and decrypt the data itself. Automated key exchange can be used to establish a temporary tunnel between two storage devices. Processes and controls are needed for the exchange, storage, authentication, and update of encryption keys. The longer encrypted data is retained, the longer the keys must be maintained and protected, and the higher the risk of their loss or compromise. Encryption standards have continued to evolve to address these risks. For example, early Data Encryption Standard (DES) key lengths of 56 bits are now considered weak because brute-force testing of the entire key space is relatively quick and inexpensive using easily available processing capacity. 3DES is considered more credible because of its effective key length of 112 bits. More recently, the U.S. government's National Institute of Standards and Technology (NIST) has suggested the use of the AES (Advanced Encryption Standard) algorithm, which employs 128-bit/256-bit block data encryption.

Longer keys are not necessarily better, however. Increased complexity in encryption is likely to bring increased latency, adversely impacting an organization's data access. As data enters the storage security appliance, it is encrypted and then forwarded. The amount of processing required to execute a security policy between storage source and destination can greatly impact what type of encryption will be acceptable. Using encryption at any level can place an enormous processing burden on the host system and degrade response time. If enhanced security requires upgrading or replacing servers, switches, routers, and arrays, then such defenses may become very costly. Dedicated storage security appliances, while adding additional initial expense, may be appropriate in complex environments with high-speed demands. Ideally, the best solution results in low levels of latency and the greatest transparency in operation at an acceptable cost.

Most enterprises leave their systems vulnerable to attack regardless of their use of encryption.

Processes and controls are needed for the exchange, storage, authentication, and update of encryption keys.

Encryption at any level can place an enormous processing burden on the host system and degrade response time.

5 Conclusion

Most organizations don't evaluate the balance of data value and security risk — risk being either a negative event or a missed opportunity — in making storage security decisions. They often fall into the trap of trying to solve the latest incident or threat without regard to its likely impact on the organization. It is easy to lose sight of the fact that storage security decisions should be made on the basis of real risk, which is the product of the probability of a threat actually occurring and the impact of that threat. A discipline of managing overall risks, rather than reacting to each threat in isolation, is fundamental to cost-effective success.

Good storage security reduces threats to data integrity while increasing the opportunity to use data to its best advantage. Qualifying and quantifying the business value of managing data threats, as well as opportunities, is critical to organizational decision-making.

Establishing security goals and determining the required pace of improvement also will provide guidance for making tactical decisions about capabilities. Without such a foundation, improvement efforts are likely to be uncoordinated and will often result in wasted resources and “reinventing the wheel.” There are extremely robust and complex point solutions that can be implemented for nearly every facet of security. Some organizations' potential risks justify taking efforts to this level, but this aggressive approach is certainly not appropriate for all organizations, even long term. These are not easy decisions.

Throwing new security tools and processes at a problem without an overall plan is likely to make the current situation worse. At best, it may frustrate users and retard the legitimate use of data. More likely, it will cause or exacerbate an organization's inability to understand its own overall security situation.

An overall approach should consider the following steps:

- Assess the risks to enterprise information in the context of business requirements and legal obligations
- Assess the security program against a reference model
- Assess security processes against a maturity model
- Incorporate security processes into overall IT and business processes
- Assess whether current data users' needs are being met
- Assess the impact of contemplated actions on the data users
- Follow change management best practices, including communication, involvement of constituencies, and ongoing feedback
- Measure, plan, build, run, and repeat

Example standards and models for reference include:

- International Organization for Standardization (ISO) 17799
- Control Objectives for Information and Related Technology (CoBiT)

Using these tools as guides and focusing on the long-term productivity of users will yield strategic advantage through superior ILM capabilities.

