

Evaluating a Storage Security Solution

**Considerations and Best Practices for
Securing Sensitive Data**

November 2006

Evaluating a Storage Security Solution

- Introduction..... 3
- Why Should You Encrypt Your Data? 3
- How Much Would a Breach Cost Your Company?..... 4
- Vendor Evaluation 4
- Technology Solutions 5
 - Infrastructure Support 5
 - Performance 6
 - Cost 6
 - High-Availability 6
 - Day-to-Day Management Overhead 7
 - Role-Based Administration..... 7
 - Encryption Strength 7
 - Encryption Placement 7
 - Auditing and Logging 9
 - Certifications..... 9
 - Vendor Interoperability..... 10
 - Long-term Key Management..... 10
 - Obtrusiveness to Infrastructure 10
 - Obtrusiveness to Your Processes 10
- Summary 11
- Appendix A – Other considerations when investigating Storage Security..... 12
- Your Data..... 12
 - What Portion of Your Data is Sensitive? 12
 - How Long do You Keep Your Data? 13
 - What Rules and Regulations Must You Comply With? 13
- Your Processes..... 14
 - What Processes are in Place Today?..... 14
- Your People 14
- Appendix B– Storage Security Solution Comparison Checklist 15

Introduction

The purpose of this document is to provide guidance into some of the factors you should consider when evaluating storage security technology and solutions. As with any security project, acquiring technology is not the only step to properly protecting your data. Part of this process should include an evaluation of the current processes and security controls in place, such as physical access controls, environmental controls, and administrative controls.

While there is no single set of requirements that applies to all organizations, this document can provide some baseline considerations.

Why Should You Encrypt Your Data?

Data in networked storage environments is significantly more vulnerable to unauthorized access, theft, or misuse than data stored in more traditional, direct-attached storage. Aggregated storage is not designed to compartmentalize the data it contains, and data from different departments or divisions becomes co-mingled in the network. Data backup, off-site mirroring, and other data replication techniques increase the risk of unauthorized access from people both inside and outside the enterprise. Partner access through firewalls and other legitimate business needs can also create undesirable security risks, and current research indicates a significant percentage of attacks come from within the firewall. With storage networks, a single security breach can threaten the data assets of an entire organization.

Data in cleartext is vulnerable to attacks. Curious or malicious insiders, administrators, partners, hackers, contractors, or outsourced service providers can all gain access to data quite easily. Technologies such as firewalls, Intrusion Prevention Systems (IPS), and Virtual Private Networks (VPN) seek to secure data assets by protecting the perimeter of the network. SAN security features such as LUN Masking and Zoning, as well as NAS security features such as access controls also attempt to address concerns about security. Unfortunately, these targeted approaches do not adequately secure storage, as data is still stored in cleartext, dangerously open to a wide range of internal and external attacks.

Encrypting your data at rest, on tape and disk, will significantly mitigate these threats and allow you to secure your data while maintaining your current service levels for operations.

How Much Would a Breach Cost Your Company?

As you analyze your protection strategies, you'll also need to consider what a breach could cost your organization. Ultimately, the value you assign to this will depend on your business and the legislation that applies to your organization. Recent studies by Gartner and other research organizations estimate the cost of a breach at roughly \$90 per customer record compromised, including notification costs, credit reporting services and administrative time.

There are many other costs that may apply, depending on the nature of your business:

- Criminal or civil penalties enforced by the courts
- Legal costs required to defend the company in such cases
- Loss of trade secrets or other intellectual property made public or falling into competitors hands
- Brand damage
- Loss of customers, or at a minimum, customer trust

Traditionally, decisions regarding the amount of security for data were based on a pretty simple assessment -- if the cost to my adversary to breach the data was higher than the data's worth, then it didn't make sense to apply the protection. However, with the myriad legislation that requires organizations to protect their customer data, this dynamic has changed drastically. The actual value of the data itself could be dwarfed by the cost of penalties enforced upon your organization. In general, identifying the general cost of a breach can help you justify a reasonable budget to put defensive measures in place to prevent it.

Vendor Evaluation

There are a range of different approaches to securing data at rest. When evaluating storage security solutions, especially those that incorporate encryption, there are some unique criteria that should be considered. For example, you may determine that you need to keep a piece of encrypted data for 10, 20 or more years. In this case, you should be very comfortable that the solution provider, and their technology, will be available to you that far in the future.

Some of the general factors you should consider when evaluating a vendor include:

- Financial stability and long-term viability
- Leadership position in the market
- Industry awards and other recognition
- Role in standards bodies, specifically those relating to encryption
- Industry partnerships
- Customer references
- Appropriate level of support available:
 - Warranty
 - Hardware replacement
 - 7/24 call center support
 - Training and Installation Support

Technology Solutions

When you begin looking at the technology itself, the primary areas you should investigate include:

Infrastructure Support

You need to ensure that the solution you choose will integrate into your infrastructure both today and tomorrow. As an example, many companies are currently concerned about data that travels offsite on backup tapes or for disaster recovery. However, it is likely that sensitive data resides in a number of other storage environments, including databases, email servers, and unstructured data in file shares. If the encryption solution does not address each of these environments, it will be necessary to manage multiple encryption and key management systems.

Questions to consider include:

- Will the security system work with equipment from your existing storage vendors?
- Does the system provide flexibility to work with other storage vendors, should your needs change in the future?
- Does the system require purchasing new storage hardware to get the encryption functionality?
- Does the system support security for data wherever it is stored (ie., NAS, SAN, DAS, iSCSI, tape, mainframe, etc.)

Performance

Storage networks offer value because they make data more accessible and more easily managed. These networks are painstakingly tuned to meet the performance objectives of users and applications. Consequently, there are significant ramifications if you add a security solution that has anything more than a minimal affect on your networks performance. Encryption solutions that are not optimized for data in storage networks can introduce unacceptable penalties: online applications may not respond quickly enough to maintain a high level of customer satisfaction; backup windows may expand to the point that you are no longer able to perform required backup operations in the time available.

Bandwidth is a measure of the amount of data that can be encrypted in a given time. Think of it as a measurement of the diameter of “pipe” for data. Look for security solutions that operate at “wire speed” -- in other words, they are able to encrypt data as fast as the links attached to it are able to move data. For applications where large amounts of data is being moved, such as tape backup or processing large files, bandwidth is the more important aspect of performance you want to evaluate.

Cost

As with any technology acquisition, you should consider the cost of hardware, software, services (installation, training, support, etc) that will require your investment not just day one, but also during the lifetime of the solution. Cost should simply be a factor in your evaluation, no more or less important than any other factor.

A big component of cost is utilization – how much of the available bandwidth of a given solution can be utilized. Or put another way, how many LUNs/volumes can you encrypt with a given piece of hardware? When you’re not encrypting everything on an array, or using the max bandwidth of a storage port, fabric-attached appliances provide a cost-effective solution.

High-Availability

As with performance, you have made a significant investment in your infrastructure to ensure that you will be able to access your data when you need to. Carefully evaluate your options for security solutions and assess how they affect this aspect of your operations. Criteria should include:

- No single point of failure, either from the security solutions themselves, or from failures in your network causing the security processing to fail
- The solution should support your disaster recovery and business continuity programs, to ensure data is recoverable wherever you need it.
- The solution should be highly available and reliable, and should not include low MTBF / high failure rate components, such as non-redundant hard drives.

Day-to-Day Management Overhead

Look for solutions that, once configured and initialized, will run with minimal interaction from your operations staff. Make sure to consider how the security solution will react if changes are made to the surrounding infrastructure. For example, is the security system impacted if you need to upgrade your host operating systems? How difficult is it to update the security system itself? Will upgrades require taking data offline? It is also important for your security solution to interface with your existing management products so that error messages and security violations are reported into the primary management system ensure timely follow-up and resolution.

Role-Based Administration

Role separation can be a powerful tool to defend sensitive data against attacks by insiders (and, to prove to auditors that these controls are in place.) The fundamental theory behind role separation is to prevent any single individual from having enough privileges to compromise all your data. One feature you might want to look for in a data security solution is the ability to define different administrator roles, which allows you to make the responsibilities of an individual as broad, or narrow, as you deem necessary.

Encryption Strength

As encryption algorithms age, and processor power increases, today's algorithms will progressively become more vulnerable to breaking. Encryption algorithms such as DES, 3DES and hashing algorithms such as MD5 and SHA-1 are generally considered to not be secure any longer. Look for products that use the strongest commercially available algorithms such as AES-256. It is also important to consider the end-to-end security of a system – encryption is only as strong as its weakest link. If you encrypt using AES-256, but store your keys in cleartext and leave them in an open operating system, it will be fairly easy to compromise the system. Because of the changing nature of encryption standards, it's also important that your encryption solution can be upgraded to address emerging standards, without requiring full hardware replacement.

Encryption Placement

Today there are several places in your network where encryption solutions can be deployed:

- Applications,
- Servers/hosts
- Storage Network
- Storage Devices

There is not a location that is best for all applications and environments - each location has its own strengths and weaknesses, and each implementation has its own nuances. Generally speaking however:

Application: When encryption is done in the application it has the most significant impact on performance. This is due to the fact that the processors that have been designed to run the application are typically the most ill-suited for encryption processing. Application encryption is also specific to a given application. If you have multiple applications that require access to encrypted data it will be difficult, if not impossible, to find compatible solutions that will utilize a common key management infrastructure. Further, it is likely that one or more of your applications will not natively support its own encryption mechanism.

Servers/Hosts: Performing encryption in your servers can be a double-edged sword. It does allow you to provision your encryption processing where it is needed. The downside, however, is that it is very intrusive to the operations of that server. If encryption is done in software, performance on that server will be significantly impacted whenever a non-trivial amount of data needs to be encrypted. If encryption is to be done in specialized hardware, there will be downtime for each server to be shutdown, have the coprocessor installed, reboot, install appropriate driver software (and perhaps reboot again), test your applications, and bring back on line. If you have tens, hundreds, or thousands of servers this will be extremely invasive to your operations. Further, as this deployment will not happen instantaneously, you will need to plan very carefully the rollout as there will be periods when some servers are encrypting data and others are unable to access it. Finally, it is important to find a solution that supports all the host configurations (both hardware and Operating System) that you have today and will use in the future.

Storage Network: Deploying encryption solutions that reside inside the storage network itself is probably the most common method in use today. The advantages of this model are that, in many cases, the solution can be deployed with literally zero downtime to your applications. Hosts and servers will not need to have their hardware nor software configurations modified, and the storage media devices themselves will not need to change. Some forethought is required to ensure you maintain accessibility to your data as well as whatever failover or other high-availability features you use. The downside of positioning your encryption in the network is that it must be compatible with whatever storage networking technologies you use, so look for solutions that can natively encrypt all open storage protocols that may be currently deployed or considered for future deployment in your infrastructure, including Fibre Channel, iSCSI and NAS protocols (NFS and CIFS).

Storage Media Devices: While not currently available at the time of this writing, there is definitely a groundswell of development underway from manufacturers of disk and tape drives. Until solutions are actually available, it is difficult to speak to the specific positive and negatives of this approach. Intuitively, it may seem that this is a very logical place to deploy encryption technology, but you will need to be confident that these solutions are able to meet your performance requirements with the limited processing power in these devices, and ensure that they utilize a key management infrastructure that will allow you to access your data today, or far into the future.

Auditing and Logging

Strong logging capabilities are another critical tool in your arsenal. You can use logging tools proactively to monitor your network looking for suspicious behavior, or reactively to investigate the events leading up to a breach.

A distinct advantage of encryption appliances is that they can simplify your area of concern. If sensitive data is encrypted, the only way to access it will be via the encryption appliance, which holds the keys to decrypt the data. Any access to data that bypasses the encryption appliance will yield only encrypted garbage. Consequently, encryption appliances can maintain very granular logs to track access to data.

However, the value these logs provide is directly related to the integrity of the logs themselves. For example, if an insider was able to perform an unauthorized function, and modify the logs to remove any trace of the activity, the log is almost useless. Look for solutions that provide extra protections on the logs – one of the best ways to do this is to have the system cryptographically sign the audit log entries. If an entry is removed, or otherwise modified, there will be an indication that a change has been made, and even the savviest adversary will not be able to entirely cover his tracks.

Certifications

It is also important to look for products that have gone through formal, independent certification. The standard certification body for encryption technologies is the National Institute of Standards and Technology (NIST), who tests and certifies third-party products against a standard called the Federal Information Processing Standard (FIPS). Other certifications, most notably the international Common Criteria standard, are also used to validate that encryption products have been built properly. Without validation, there is no way to ensure the products perform as promised. With encryption, this is even more critical.

Vendor Interoperability

When you decide upon a storage security solution it will be required to work with your storage infrastructure both today and tomorrow. You should give significant weight to solutions that have gone through interoperability testing with leading storage vendors. Disk and Tape providers as well as backup application vendors, network solution providers, and possibly even your specific applications should be considered. A storage security company's willingness to complete these steps should reduce any concerns you might have about their ability to make their products work in your environment.

Long-term Key Management

The key management system may well be the single most important component of your storage security solution. As we've discussed, you will likely need to maintain keys for many years. You need assurance that the keys used to encrypt data will be available whenever and wherever authorized access to data is required. At the same time, the keys need to be secured so that they themselves aren't compromised (resulting in a data breach). However, key management systems come in many shapes and sizes. Criteria to consider are:

- Does the key management system store the keys securely, giving you confidence that key will not be accessed without authorization?
- Does it allow you to securely replicate keys wherever you need them to be?
- Is it "storage aware"? Does it allow for functions like key expiry, making data read-only after a certain date, destruction of data with key deletion, etc?
- Will the key management system interface with other security components as they are added to your network?
- Does the vendor have a roadmap that you feel confident will give you the features you need in the future, while guaranteeing accessibility to any key material you use now?

Obtrusiveness to Infrastructure

Look for solutions that easily fit into your storage infrastructure. During the deployment phase you will want to keep interruptions to your network to a minimum, with zero-downtime as the ultimate goal. Look for solutions that do not require you to take hosts, switches, or other devices offline to deploy. Some solutions will require you to take data offline during an initial encryption, or subsequent rekeying phase, or expose the data to a prolonged window of vulnerability while data is processed via a scripted host based solution. Look for a solution that can perform these functions transparently (in the background), with the data fully available to whoever may require it.

Obtrusiveness to Your Processes

You have already spent a great deal of effort trying to maximize the efficiency of your people and systems. If a new technology solution causes you to drastically modify the way tasks are completed, it is unlikely to be a positive change. However, a properly

engineered storage security solution may actually increase the efficiency of certain tasks. For example, encrypting backup tapes may allow you to more quickly and cost-effectively move them to your disaster recovery site.

Summary

There is no single “magic bullet” that will solve your information security requirements. A comprehensive storage security plan will include inventorying your current information, assessing your current processes technology, and evaluating technology solutions. Technology assessments should include consideration of the performance, cost, key management, certifications, and vendor pedigree to ensure that the solution you chose will meet your requirements today and in the future.

Appendix A – Other considerations when investigating Storage Security

Your Data

What Portion of Your Data is Sensitive?

One of the first steps you should undertake is a general inventory of your data. This may be something you have done before, but it is a worthwhile task to consider again. You should look at the data with an eye towards how much time and energy should be spent maintaining its confidentiality, integrity, and availability.

Confidentiality is a measure of how safe the data is from being viewed by unauthorized personnel. The percentage of data that requires strong confidentiality protections will vary greatly from organization to organization. At a minimum, you should protect information that your customers and employees would consider private, as they would potentially suffer harm if that information were lost, stolen or disclosed. Properly implemented, storage security solutions can significantly improve data privacy.

Integrity is, in a sense, a measure of how “good” the data is. Has it been tampered with? Was it entered into the system properly? Will it remain stable over a period of many years if it needs to be retained? While, in general, storage security solutions cannot guarantee data was entered properly, they can do a great deal to help ensure that it hasn’t been improperly tampered with.

Availability is a measure of how quickly and easily you are able to access data. For example, if you have a piece of data that needs to be accessed several times a day, it wouldn’t be very efficient to keep that data stored only on backup tapes in a secure location. Ideally, storage security solutions should have little or no effect on the availability of your data.

It is worth noting that as security solutions have become more affordable and less obtrusive; many organizations will simply secure all their data. This strategy doesn’t eliminate the need to inventory your data, but it can make the process of determining what data should be protected much easier – simply secure everything.

In the context of evaluating storage security solutions, the primary reason to inventory data is to get an idea of the magnitude of data you will be protecting. The breadth of resources involved (hosts/servers, networks, disks, tapes, etc) will play a factor in determining the best solution for your organization. If, for example, you determine you only need to encrypt one backup tape a week (an unlikely scenario) then perhaps a software-based backup utility will meet your requirements. For larger enterprise

customers, the amount of data is likely to mandate the performance of a hardware-based solution.

How Long do You Keep Your Data?

Retention times are an aspect of availability, and they can have a significant impact on your solution choice. Some organizations have retention requirements that compel them to keep data on hand for a few years. However, as retention legislation becomes more common, it is not unprecedented to find organizations that have to keep data for decades, if not forever. This creates some logistical challenges, such as where to store the physical media, or how to read the media back in the future (e.g., will there be a tape drive that can read xxxy tapes?)

If you are considering storage encryption solutions, retention times will drastically affect the number of encryption keys you need to maintain. You will want to carefully consider solutions that will both hold a great number of keys, as well as restore them in a timely fashion.

What Rules and Regulations Must You Comply With?

It may be valuable to ask: why are you considering a storage security solution? Some organizations will deploy encryption simply to ensure compliance with various laws and have no independent desire to add extra protection to their data. Other organizations want to protect both themselves and their data from any legal requirements. Most organizations fall somewhere in the middle. Some of the more common legislations driving storage security deployment include:

California SB 1386 (and similar laws in 20+ states) - California's Senate Bill 1386 went into effect July 1, 2003, with significant implications for any organization that does business in California. SB1386 seeks to reduce identity theft and protect California residents' right to privacy by requiring disclosure of any breach to the security of a computing system where there is a reasonable belief that an unauthorized person has acquired unencrypted personal information.

The Gramm-Leach-Bliley Act (GLBA), which took effect July 1, 2001, ushered in sweeping changes for the financial services industry. Along with enabling new types of financial services for consumers, the act also requires institutions to develop extensive administrative, physical, and technical safeguards to secure the confidentiality of customer records and information.

CISP - In April 2000, Visa launched its Cardholder Information Security Program - a very specific set of mandates designed to protect its cardholders from identity theft and other misuse. Visa outlined key security requirements, along with a program for validation and auditing. In December of 2004, Visa and MasterCard joined forces to simplify compliance for merchants and payment processors with the jointly-developed, 12 point PCI standard. The scope of these requirements is

quite broad, incorporating best practices for perimeter security, data privacy, and layered security.

The Health Insurance Portability and Accountability Act (HIPAA), which took effect in April 2003, has been a primary focus for the healthcare industry. HIPAA's Security Rule mandates a wide range of administrative, physical, and technical safeguards to secure the confidentiality of patient records. HIPAA regulations require business and IT managers to secure electronic protected health information (EPHI) from a broad range of potential threats. The stakes are high: penalties for noncompliance include criminal prosecution, fines, and up to 10 years in prison.

Federal Senate Bill S.1789 *"A bill to prevent and mitigate identity theft, to ensure privacy, to provide notice of security breaches, and to enhance criminal penalties, law enforcement assistance, and other protections against security breaches, fraudulent access, and misuse of personally identifiable information"* is currently going through the process to become a federal law. The goal of S.1789 is to protect consumer's identities and private information from the many threats that exist today. The bill is broken into sections that detail the status, purpose, applicability, requirements and penalties for non-compliance among other topics.

The bill places significant responsibility on business entities and data brokers to provide security controls that are adequate for the sensitivity of the data involved. Furthermore, the business carries the responsibility to notify all consumers whose information was subject to a security breach. Any business involved in interstate commerce that handles privacy information on more than 10,000 United States persons falls under the requirements laid out in the bill.

Your Processes

What Processes are in Place Today?

It is important that you also analyze the processes you use today, and evaluate how they protect you against the risks you face. For example, if you never move data, disks or tapes outside of a single computer room, protected with strong physical security, and are confident in the personnel you have maintaining and enforcing these processes, the requirement for encrypting data is probably less than that of organizations with Disaster Recovery facilities, move backup tapes to an offsite location, etc.

Your People

Whichever solutions you end up deploying, do not lose sight of the importance of properly maintaining your hiring processes, and training your employees to properly utilize the security features of your solution. Even the most secure technology will not properly protect your data if it hasn't been installed, configured and managed properly.

Appendix B– Storage Security Solution Comparison Checklist

Evaluating a Data Storage Security Solution			
Selection Criteria (yes/no)	Vendor 1	Vendor 2	Vendor 3
VENDOR PEDIGREE			
Relevant number of years in business			
Financial stability and backing			
Relevant industry partnerships			
Customer references			
Role in standards bodies, specifically those relative to encryption and key management			
Total YES			
SERVICES AND SUPPORT			
Strong pre-sales process to support assessment and proof of concept			
On-site installation and configuration			
Comprehensive professional services			
On-site technical support if required			
24x7 global technical support			
Strong pre-sales process to support assessment and proof of concept			
Comprehensive training			
Total YES			
INFRASTRUCTURE SUPPORT			
Breadth			
Works with existing storage equipment			
Flexibility to support future storage needs			
Unified platform to secure data in all storage environments (NAS, DAS, IP-SAN, FC-SAN, mainframe and open systems tape)			
NAS/DAS/IP-SAN Support			
Supports CIFS/NFS/iSCSI			
Supports IPsec link encryption for end-to-end Security			
Integrates with and enforces directory services (Active Directory, NIS, LDAP)			

Supports in-place encryption of existing data with zero downtime			
SAN Support			
Can be deployed transparently in-line, or fabric-attached			
Zero downtime deployment			
Disk and tape encryption supported from the same solution			
Supports host authentication to strengthen fabric security			
Supports in-place encryption or re-keying without taking data offline, at 20MBps or higher speeds			
Can compartmentalize shared storage by encrypting each LUN with a different key			
Tested and interoperable with multi-pathing software for failover and availability			
Tape Support			
Employs hardware-based compression before encryption to ensure media usage is not increased			
Flexible encryption granularity to support key per tape, key per tape pool, or key per host			
Integrates with relevant backup software			
Total YES			
PERFORMANCE			
Can be deployed without degrading host or server performance			
Data throughput meets storage I/O Requirements			
Total YES			
TRANSPARENCY			
Can be deployed without taking existing data Offline			
Can be deployed without installing software on clients or hosts			
Total YES			
HIGH AVAILABILITY			
If hard drives (which have a high MTBF) are			

used, they are redundant to prevent a single point of failure			
Solution can be clustered for high-availability			
Cluster members be added/removed while the cluster is online			
Supports disaster recovery programs, ensuring data can be decrypted where it's needed			
Total YES			
ENCRYPTION			
System uses the strongest commercially-available encryption algorithms (such as AES-256)			
Encryption is optimized for storage to prevent chosen text or dictionary attacks			
A True Random Number Generator (TRNG) is used to create high-entropy for strong keys			
Alternative versions of decryption available in the event that primary encryption/decryption resource is unavailable			
System is engineered to enable hardware/software updates to address emerging standards			
Total YES			
KEY MANAGEMENT			
Key management is fully automated, so keys can be distributed and archived without intervention			
Encryption keys are never exposed in cleartext			
Keys can be securely replicated to other locations for DR or information sharing			
The system allows for policy-based expiration of encryption keys in accordance with data retention policies			
The key management system can support key generation and management for third party encryption systems			
The solution provider has a solid roadmap for next generation key management capabilities			
Total YES			

ACCESS CONTROLS			
Role-Based Access Controls (RBAC) are available to support role separation for administrators			
System offers two-factor authentication for administrators			
Total YES			
ADMINISTRATION AND MANAGEMENT			
Minimal impact on day to day operations			
Administrators have enterprise-wide control, even for remote systems			
Provides an intuitive management interface, and also support CLI for scripting of common tasks			
Interoperates with industry-standard protocols like SNMP			
Total YES			
LOGGING AND AUDITING			
Configurable to track desired events and administrative actions			
Audit logs are cryptographically signed and tamper-proof			
Logs can be exported to Syslog or other reporting tools			
Total YES			