Storage Security Best Practices

GlassHouse Whitepaper

In a recent GlassHouse survey of storage management, 52% of respondents graded their storage department "Fair" for security readiness.

Storage must become an integral part of the corporate security strategy

The widespread adoption of networked-based storage, particularly Storage Area Networks (SANs) has provided many benefits including improved utilization and better access to data within corporate environments. However, this transition from captive storage tethered to servers to today's independent networked devices has largely taken place under the radar of corporate security groups whose traditional focus has been on areas such as intrusion detection/prevention and protection of host systems. As a result, the storage infrastructure – both primary storage and especially secondary data storage like backups and archives – is likely to be a company's Achilles' heel when it comes to security. Policies for data security are a corporate concern and should be integrated as a fundamental element of an enterprise security strategy. Strategic security policies can then spawn tactical and operational policies through the joint efforts of the security and storage organizations.

This paper details several key steps that should be taken to ensure that information both in transit and at rest is appropriately maintained with regard to confidentiality, integrity, and availability. The recommended approach combines extension of standard security best practices into the storage realm with particular focus on factors unique to protecting off-site storage.

Adopt a multi-layered security approach

Many organizations have evolved to a high degree of maturity with regard to network security. When it comes to security, a SAN is no different from any other network. The switches, systems and other devices that make up the corporate SAN must be protected using the same or similar approaches as those used to secure the corporate LAN. This requires a multi-layered approach including:

- **Data classification** – Establishment of policies for identifying and classifying corporate data with regard to sensitivity of information is an important first step. Corporate intellectual property, personal information about customers or employees, and other sensitive information that is sitting on storage media throughout the environment. Knowing where this data resides enables the focusing of resources appropriate to the level of protection needed.
- **Authentication** – Just as IP networks have addresses that if not properly protected can be breeched, fiber channel SAN devices must be similarly protected. SAN World Wide Name (WWN) addresses are subject to threats such as spoofing, and should be protected accordingly. Widespread use of techniques like WWN zoning expose a storage network to authentication breeches that can be avoided without significant cost. Functions such as port-based zoning, port access control lists (ACLs), and multi-level authentication are readily available in fibre channel switches.

In a recent GlassHouse survey of storage management, 54% of respondents indicated that they have no documented security procedures in place for their storage infrastructure.

- **Authorization** – A common practice within both storage management and backup management is to provide administrators with full "root" access or its equivalent. Rather than to enforce privileges based on role and responsibility. Many SAN switches, storage management applications, and enterprise backup products provide the capability for role-based administration and this functionality should be leveraged where possible.
- **Encryption** – All data deemed to contain sensitive information should be encrypted when it is stored or copied. In addition all management interface data transmitted over any non-private network should be encrypted. The challenge of encrypting data at rest will be discussed in more detail below.
- **Auditing** – Logs of administrative operation by any user should be maintained to ensure traceability and accountability. This information is also critical to demonstrate compliance with corporate data management policies. Wherever possible this information should be rolled up into a storage or data management reporting tool.
- **Segmentation of duties** – where data is super sensitive it may be prudent to ensure that the person authorizing access is not the person charged with responsibility for execution.

Duplicate your removable media

Depending upon a single copy of data is never a good idea. While tape media can have a long life, it is susceptible to environmental and physical damage. A common practice is to perform nightly backups and then ship those tapes offsite – with no verification process. The recommended best practice is to duplicate or clone backup tapes and then to send the copy offsite. It is important that this process be executed by reading the original tape and writing a copy of it rather than simply using an OS feature to create two similar and possibly faulty tapes. This has the benefit of both verifying that the backup data is readable as well as eliminating the single point of failure of the backup tape.

The reason most often given for not having of a duplication policy is lack of time. From a practical standpoint, backups take too long, and it is impossible to duplicate the data in a timely fashion. There are a number of ways to address this issue. First, the original copy takes place within the designated backup window. The second copy can then be made "off line" so to speak, or outside the original backup window that disciplined the time available to write the original copy. With today's high performance tape drives, the ability of many backup products to stream to multiple devices simultaneously, and especially the increasing popularity of disk-based devices, like virtual tape libraries (VTLs), there is no acceptable excuse for not maintaining a redundant copy of backup data.

Ensure media management policies are in place and followed

Media management refers to the act, manner, handling, supervision, or control of media (usually, but not always, tape). The ultimate goal of successful media management is to preserve the integrity of the recordable resources. The following items should be considered with regards to media management.

Removable media should be tracked by bar codes and reports should be generated on a daily basis detailing the current location of the media. A best practice is to report daily on tapes that are to be sent offsite and those that have expired and should be retrieved from offsite storage to be recycled or destroyed. Documented standard operating procedures should be in place to ensure that these are carried out.

The offsite location and the process which is used to access the offsite storage should be analyzed for security concerns. Media should be placed in locked tubs before leaving the data center and subsequent tracking done at the "tub" level. Tubs of media should be signed for and should never be left exposed for someone to pick up.

Reconcile the inventory of media that is stored offsite on a monthly basis. At the end of each month, a physical scan of the offsite storage should be compared to the records of the backup/archive application to discover inconsistencies. If media is not accounted for, then appropriate steps must be taken.

The latest cartridges contain mechanisms that are not exactly hardened against misuse. A dropped tape cartridge is probably a cartridge that won't subsequently work. It would be great if media manufacturers could harden their cartridge design, or at least include a chemical indicator of impact but until then, this issue can only be addressed through training and continuous reinforcement. How many times have you seen an operator with an armful of cartridges have two or three or all of them hit the floor?

Once the media has reached it obsolescence or can no longer be relied upon for its integrity then the media must be appropriately destroyed. There are services that will buy recordable media and return a certificate of destruction. The destruction of magnetic media is usually accomplished by applying a strong magnet to the cartridge, thus scrambling the data on the tape, rendering it useless. DOD standards call for triple erasure to be absolutely sure.

Dust is the enemy of most media and media recording devices. To a magnetic head, a particle of dust is like a boulder on the freeway to an individual. The backup and archive environment must remain clean and dust free. A soft static-free cloth should be used to clean the outside of cartridges and dust should be removed from slots of a library or storage rack by using compressed air from a spray can. Tapes should be shipped in an electrostatic holder in a tub and not piled into a tub or a cardboard box. Although appearing quite durable, tapes can be easily damaged by being mishandled. Another critical element in secure media handling is to ensure that your offsite storage vendor also follows best practices. Here is a quick list of some of the things you should consider.

In a recent GlassHouse survey of storage management, 79% of respondents said that they were not encrypting their data.

- **On-Site Vulnerability -** Don't leave your tapes in an unlocked container (like an open cardboard box) on the receptionist's desk to await pickup. Pick up should be a standard operating procedure where a responsible IT person hands over and receives a signature from a known (ID carrying) vendor representative.
- **Background Checks -** The company is going to be storing all of your critical data, so you must trust every one of their employees. Therefore, it is critical that they are performing background checks on those employees
- **Barcode Support** - The company should be able to understand barcodes on either your tapes or tubs/cases. This will allow them to scan in the tubs/tapes before they leave your facility, and scan them into the vault once they arrive back at their facility. This will help a lot when you are trying to find a tape that was lost in the process.
- **Secure Transport** - You should talk to them about the entire process of how your media is handled from start to finish. Look for an emphasis on physical security, and audit and control mechanisms to ensure that the process is being followed. It is inadvisable to move sensitive data from point A to point B in truck or vehicle emblazoned with the vendor's name, easily identifying it as carrying sensitive data.
- **Container Vaulting** - Container vaulting is when you send them a tub/box of tapes, and they track only the tub/box. They should support this type of vaulting.
- **Individual Media Vaulting** - Individual media vaulting is when you send them a tub/box of tapes, and they track each piece of media in the tub/box. They should support this type of vaulting.
- **Physical Security Controls** - It should be impossible for you to get in to the vault. You should be able to get access to your media with proper identification, of course, but you should never be able to just go in the vault.
- **Environmental Controls** - Tapes and other media should not be stored in a vehicle's trunk, or any other non-environmentally controlled location. If they are going to be storing your tapes, the environment must be strictly controlled, including temperature, humidity, and static control.

Encrypt all data that contains sensitive information

The final and perhaps most critical layer of a multi-layer protection strategy is data encryption. If all of your other defenses have been compromised, an effective encryption capability renders the data unusable to unauthorized individuals. While encryption capabilities have been available for some time, few organizations have taken advantage of them. This is due in part to a lack of appreciation of the actual risk as well as because of technological limitations that made encryption impractical. Awareness of the risk is no longer an issue, so let us focus on the technological concerns.

As discussed previously, security of storage overall has not been a priority for most organizations, and storage administrators have traditionally focused on convenience and manageability rather than security. When considering encryption specifically, the practice of encrypting information that travels over publicly accessible networks is

standard procedure today. The use of virtual private networks (VPNs), secure web sessions for sensitive transactions through secure socket layer (SSL), and secure login sessions via secure shell (SSH) are examples of encryption in use by millions of people every day.

However, these are all examples of securing *data in transit*. The focus is on preventing someone from unauthorized snooping as data travels from point A to point B. The issue of securing *data at rest*, whether it resides on disk, tape, or optical media, is a more recent concern. Data encryption can occur at the application, database, operating system or network level. In addition, for media to be shipped offsite, backup vendors also provide encryption capabilities. Let's examine each approach.

Application/Database Encryption

For database applications, there are both native and third-party tools available that can be employed to encrypt sensitive tables, or rows or columns within tables to ensure that the data is viewable only by users specifically authorized to see those elements. Fields such as credit card numbers, employee salary information, and personal medical data are often encrypted in this way. The advantage of application-based encryption is, of course, that it is tightly coupled with the application itself. Only the data that needs to be encrypted is, thereby minimizing overhead. It also has the advantage of encrypting data at its origin, thereby providing an even greater level of security.

Operating System Encryption

Modern operating systems also incorporate the ability to encrypt data stored within a file system, directory, or individual file. For example, Windows XP provides an Encrypting File System that supports public key encryption and is fully integrated within the operating system. Encryption at this level can be very flexible and applied to a wide range of data. The potential downside is CPU overhead and file system performance as data must be decrypted to be accessed. Key management is another concern, as will be discussed later.

Network Encryption Devices

At the network-level, a new breed of appliances has emerged that can provide a transparent level of data encryption. These appliances can sit on either a SAN or LAN and provide encryption at wire speeds to either some or all of the data that travels on the network. These devices have several advantages that make them worthy of consideration.

First, they are fast. One of the major roadblocks to enabling encryption is performance (see the discussion of backup encryption below). These network devices perform encryption at wire speeds and introduce minimal latency to the environment.

Second, they are versatile. Depending on the organizations specific encryption requirements and network architecture, there is likely a device that will meet its needs. Various devices are optimized to support a variety of host systems, network protocols, primary storage, backup, and even application-specific requirements, such as database

field encryption. Residing at the network level, these devices can often be managed with little or no impact on host systems or applications.

Third, they are relatively inexpensive. Given the risk of privacy exposure or loss of proprietary information, the return on investment of these devices can be compelling.

Encrypting through backup application software

With regard to backup specifically, most major backup applications offer encryption either as a standard or optional component of their product. Let's examine how this works.

An enterprise backup application typically consists of several components:

- A central master server that maintains a catalog of all of the backup information along with the policies, schedules, configuration information, and administrative components of the environment.
- One or more servers that manage the creation of and access to backup media (either disk or tape).
- The systems whose data is actually backed up – the backup clients – who transmit data to the servers to be processed.

Traditionally data and metadata is sent in the clear between the systems involved in this process and no encryption is performed. The exception, in most environments, is usually when a remote backup service provider is used rather than an in-house operation. Because this data is sent over an external network, it is common practice for third parties who provide backup services to employ encryption.

For an in-house operation using an enterprise backup application, a typical approach to employing encryption might consist of:

- enabling encryption within the backup application
- encrypting data at the backup client, and then
- transmitting the encrypted data to the appropriate backup server on to tape or disk.

It sounds simple, so why isn't this common practice? Unfortunately there are several reasons:

First, the processing impact of performing host-based encryption on a production environment can be substantial. In an active environment, where critical business functions are being performed on a 24-hour basis, this overhead may be overly intrusive.

Second, the encryption process can slow the rate of data being transmitted to the backup host and then subsequently to tape. Modern tape devices have very high throughput rates and require a constant stream of data to perform in an acceptable manner. When the data rate drops to a point below the streaming requirements of the device, performance dramatically falls. Essentially the tape device must stop, backup, and

restart. When this happens repeatedly, "shoe shining" occurs, resulting in very poor performance and quicker tape wear.

Third, encryption of data has the side effect of "flattening" a file, i.e. making it less compressible. Modern tape devices perform compression and most environments depend on compression to minimize the number of tapes required for a nightly backup. Compression also impacts on performance as well, since compressed data represents fewer actual bytes that must be written to tape. While it is possible for backup applications to compress data at the client side, as with client-based encryption, the performance implication on production systems is usually unacceptable.

The final and possibly most critical reason that encryption has not been more widely adopted for backup data has to do with the management of encryption keys. An essential element of the process of encryption and, more to the point, decryption key management is mission critical. Restoring encrypted data is impossible without the key and having a process for administering, protecting, and recovering keys becomes an essential parallel process for data recovery.

Managing Encryption Keys

All current encryption products use a variation of private key cryptography. The loss of a key, or of the encryption product in use, would cause the loss of all encrypted data. As discussed above, encryption is critical to data management, especially long-term storage of offline data. No alarms might sound if a backup tape is purloined and read, and these tapes may contain consistent copies of entire data sets. In fact, without encryption, backup tapes are normally completely insecure. However, encryption can turn the tables on security, locking your own data away from your use.

Consider encrypted data on tape as being locked in an indestructible safe sitting out on the street. The keys to that safe are securely held in your pocket, so you can access that data if needed. Nevertheless, what happens if you need to share those keys with someone else? Once they get away from your control, they can be duplicated and your data can be read. In addition, what happens if the keys are lost?

Best Practice Key Management

The first factor to consider is the need to change the keys. If you have delegated encryption to an employee, you would want their key to become useless if they are no longer trusted. So you switch to a new key, encrypting future data with it.

But this raises a problem: What becomes of the old data? There are two possibilities: Either the old data continues to use the old key, or it is updated to use the new key. The second, though obviously preferable, is technically extremely difficult to achieve. Updating old data on disk requires re-reading and encrypting it with the new key, a time-consuming process. However, re-encrypting data on tape, especially offsite tape, is more challenging still. All old tapes would have to be recalled and rewritten.

If data is not re-encrypted, the question becomes one of how to deal with data encrypted with two different keys. If all new data uses the new key, and all old data uses the old, you will have to switch between keys based on the data set being read.

Preventing Key Loss

The question of changing keys, and the possibility of data loss, raises another consideration: protecting keys. It is essential that the right key is available to read any data set in the future. No matter how well your application deals with old and new keys, no application will be able to read encrypted data without a valid key.

One best practice for encryption keys is known as key escrow. This is a service, normally provided by a trusted third-party, which holds keys offsite in non-volatile media in case they are needed. Key escrow is akin to entrusting a lawyer with a document, or locking a key in a safe deposit box. It ensures that your key will be available to you in the future and available within a time frame that supports committed RTO. This is especially critical for disaster recovery preparation, since all on-site resources, including encryption keys, could be lost.

Some modern systems also include specialized key recovery functions. Decru, for example, allows a quorum of key recovery cards to recover lost keys from their encryption device. In this system, each trusted individual is given a special key recovery card. These cards are useless on their own, but can be used to convince the application that an encryption keys is being requested for a valid purpose. While not a full disaster recovery solution, this key recovery technique provides an additional level of protection.

Protocol Upgrade

Just as changing keys can make data unreadable, switching from one encryption protocol to another is a challenge. Remember, this is exactly the same problem you already manage as tape technology matures and we move from reel to reel, through DAT, to DLT to SDLT to LTO etc. As with a change in tape technology, so with encryption technology, either all data must be re-encrypted or a dual-protocol system must be used. The latter is not preferable, however, since the rationale for switching from old protocol still applies to that old data.

Why switch protocols? Weaknesses in encryption protocols are discovered fairly frequently, and older protocols eventually fall out of favor. You may also choose to move from one encryption vendor to another and different vendors support different protocols.

Therefore, a procedure must be developed for a protocol switch. This will be a time-consuming process, requiring a great deal of personnel and technical resources. All online and near-line data must be re-encrypted, and all offline data (backup tapes, DR copies, and archives) must be recalled. Some vendors provide protocol upgrade tools, but these merely assist in the mechanical re-encryption procedure, not the overhead of implementing this process.

Preparing for Loss of an Encryption Vendor

One of the biggest crises for encryption users would be the loss of the vendor of their encryption technology. Although it would presumably still function after the vendor shut their doors or cancelled the product, it would be foolish to continue using it.

For this reason, the loss of an encryption vendor should be considered similarly to a protocol upgrade. That is, all old data should be recalled and re-encrypted with a new scheme. The large cost of this move should be considered when evaluating products. Which vendor is most-likely to continue supporting their product, even if it is no longer current? Some vendors might even be willing to post a bond or place their software code in escrow to protect against this scenario.

Getting Started with Encryption

It should be clear that implementing data encryption requires planning and preparation. Getting started begins with developing the strategic policies concerning what data needs to be encrypted and then identifying that data and any copies of that data within the enterprise storage environment.

The next step is selecting an encryption method. The type and quantity of data to be encrypted, the capabilities of the existing reference architecture, constraints imposed by time windows, physical logistics, and RTO requirements must all be considered. For environments where limited amounts of data are to be encrypted, then application/database encryption or backup application level encryption may be appropriate. For encryption of large amounts of data on selected hosts, file system encryption might be a good choice. For wide scale encryption needs, then a network-based encryption appliance would be the most likely choice.

With any approach, the management process to secure the encryption key needs to be addressed. The key needs to be totally secured yet readily available in time of need, and readily returned after the need has passed. Some of the appliances provide assistance in this area making them even more attractive. Operating system vendors also provide guidance in this process, as well.

The standard operating procedures (SOPs) governing security of data at rest must contain a metrics base that tracks not only completion and compliance but also the logistics management of both the physical data container and most importantly, the encryption key itself.

Conclusion

Securing data, both in flight and at rest, is a cross-functional responsibility. Storage security needs to be fully integrated into the corporate security framework. As a first step, a complete review of storage and backup security policies and practices is recommended resulting in a plan of action. Existing network security standards should be applied to storage networks and policies and procedures specific to the unique characteristics of the storage environment established.

Finally, everyone who manages, administers, or operates IT infrastructure needs to become fully security conscious. Security is as much a culture of awareness as it is a corporate policy directive. To truly protect the organization's critical data, continuous focus on culture, practice, and control is imperative to a successful security strategy.

About GlassHouse Technologies, Inc.

GlassHouse is the leading provider of independent services that help organizations solve the business problems of enterprise storage. From strategy through implementation, operations and customer support, GlassHouse partners with clients to achieve predictability and manageability in storage operations, enabling cost control, risk mitigation and increased service levels. GlassHouse clients include UBS, Exxon Mobil, Charles Schwabb, Virgin Mobile, and The Guardian Life Insurance Company of America. More information about GlassHouse is available at www.glasshouse.com.