Securing Sensitive Information

Silverton Consulting, Inc.
StorInt™ Briefing

An examination of information security issues, methods and securing data with LTO-4 tape drive encryption

Introduction

Each month many companies, big or small, well known or unknown, experience a data security loss with the potential exposure of thousands to millions of sensitive customer or employee records. Recent regulatory actions have made such losses much more onerous. Corporations need to reduce the financial risks of a security breach as well as protect their brand reputation. As such, corporate management is looking to CIOs to minimize these risks with effective security for all sensitive corporate data, wherever it may reside.

Encryption has emerged as a best practice mechanism to security breach risk. As an important consideration for corporate officers cryptographic methods will be examined that can mitigate risks associated with data security breaches, specifically tape data encryption. LTO technology is the most widely adopted data storage tape format and as such, LTO-4 drive encryption will be thoroughly examined along with an LTO-4 encryption user case study.

Using encryption in data protection

Data security legislation aimed at protecting the disclosure or loss of personal data is now in force internationally. Such laws increase the due diligence and care businesses must take with the data in their possession. Some even incorporate penalties for businesses failing to comply. Examples of legislation concerned with data protection include:

- Sarbanes-Oxley (US)
- Gramm-Leach-Bliley Act (US)
- USA Patriot Act
- European Union Data Protection Act
- AIPA (Italy)
- GDPdU and GoBS (Germany)
- Electronic Ledger Storage Law (Japan)

Within the US alone the Privacy Rights Organization found that over 245 million data records containing sensitive personal information have been exposed to security breaches since January 2005[1]. For example,

- On February 25, 2005, a commercial bank's backup tapes were lost containing social security numbers and other personal information of over 1 million customers.

[1] All information on security breaches were found on http://www.privacyrights.org/ar/ChronDataBreaches.htm as of October 26, 2008

- On May 22, 2006, a government agency reported the theft of a laptop and computer disk drive containing the social security numbers of over 28 million American veterans discharged since 1975.
- On January 17, 2007, a large retailer reported "unauthorized intrusions" over a number of years in their customer transaction processing system potentially breaching security for over 47 million credit and debit card account numbers.

Public notifications such as these are a consequence of data "security breach notification" laws introduced in 34 U.S. states over the last two years. Breach notification laws require all companies doing business in a state that owns or licenses computer data, which includes personal information, to disclose any breach of security of that data. Most of these laws specifically exclude any notification requirement for securely encrypted data.

In addition, it is estimated that notifying clients of a data security breach could average $50 per record[2], but the real cost to a business reaches beyond fines and penalties. Data loss could cost millions in lost revenue, loss of customers, intellectual property, and damage to the brand.

Cryptography is the answer

Cryptography is defined as the study of the practice of hiding information from unauthorized view. Cryptography uses ciphers, essentially encryption and decryption algorithms, together with a key, to render data unreadable to unauthorized personnel. Encryption takes plain- or cleartext information and passes this data through an algorithm transforming it into ciphertext which appears to be random data. However, with the proper key and cipher one can decrypt the ciphertext back into cleartext and use the data normally.

Two encryption techniques have evolved to secure sensitive data including:

- **Data-in-flight encryption** - encrypts data while it is being transmitted from one place to another across a network. For example, as transaction data is encrypted for transmission across the Internet the prefix of a website changes from "HTTP" to "HTTPS". This type of encryption is only temporary. Before and after data transmission, the data exists as cleartext, readable at either end, but while in transit the data is encrypted.
- **Data-at-rest encryption** – encrypts data before it is stored and decrypts the data after it is retrieved. For example, using LTO-4 data encryption, the stored data is encrypted when it is written and decrypted when an authenticated user reads it. This type of encryption is permanent and protects data where it is stored. The data is in cleartext only before it is written to the cartridge or after it is read from the cartridge.

Data-at-rest encryption protects stored data from unauthorized access and physical theft provided proper diligence is maintained in protecting encryption keys. A potential thief

[2] ***InformationWeek,*** April 11, 2007

could steal a tape cartridge but not be able to read the encrypted data stored thereby preventing a security breach and the associated costs to the organization.

As such, data-at-rest encryption has become more broadly available, and provides formidable protection of sensitive data in the case of physical theft security breaches. Vendors supply this encryption technique via several methods including:

- **Software encryption** - The operating system or application software automatically encrypts data before it is stored and decrypts data after it is read. Today Windows, Mac OS/X, Linux, Solaris, HP-UX, AIX, zOS and others support some form of host software encryption for data security. The downside of software encryption is that a lot of CPU cycles are typically consumed and can significantly burden servers generating lots of data.
- **Appliance encryption** – Network attached hardware encrypts and decrypts data flowing through the appliance as it is moved to/from disk or tape. Such products must be placed in the data path between all hosts and storage devices handling encrypted data. The upside of encryption appliances is that encryption/decryption does not execute at the host. However, the cost involved with adding these appliances can be substantial. Also, when one considers several appliances may be required to serve a multi-drive library, the management and scalability issues become apparent.
- **Switch encryption** – Similar to encryption appliances, these products encrypt and decrypt data that flows through their FC switches. Currently, Brocade supports switch encryption for disk data and Cisco supports switch encryption for tape media. The advantage of switch encryption products is that they are already in the data path and as such, are most likely highly available. On the other hand switch encryption is expensive and performance requirements of both switch and encryption needs, although related, must be determined independently.
- **Subsystem encryption** – Subsystem firmware and/or hardware encrypts and decrypts the data before it is stored on the back end disk drives. Most large storage vendors such as EMC and HDS offer this feature for their high-end, enterprise class storage. Similar to switch encryption, the advantage of subsystem encryption is that the subsystem is also already in the data path and so, most likely highly available. However, these subsystems may not scale well for encrypting large amounts of data. Also, there may be a minor storage performance penalty for using encryption.
- **Device encryption** – Data is encrypted and decrypted at the device level. The open standard LTO-4 format specification includes the ability for data to be encrypted by the tape drive hardware. Encrypting tape drives are offered by LTO tape drive technology vendors. The advantages of device encryption are that there is typically no noticeable tape drive performance degradation and it has high scalability by simply adding more drives. Moreover, device encryption does not need high availability as tape media could easily be swapped to another transport and disk data could be reconstructed from other drives in the RAID group.

Tape storage has low total cost of ownership and supplies significant benefit to organizations by providing archival storage, offsite data protection, ability to address

compliance requirements and low energy consumption. Most businesses also use tape backup to protect primary data and move backup tapes to secure offsite locations for disaster recovery purposes. Because of this mobility and sheer number, tapes are vulnerable to physical data theft. As such, most data centers should consider device data-at-rest encryption, specifically tape drive encryption, for sensitive backup tapes.

> **"Have we not learned from history yet? If you're going to give (data) to a third party then you either encrypt or password-protect it."**
>
> **Linda Foley, executive director of the Identity Theft Resource Centre San Diego**

The Encryption Process

Historically, ciphers were characterized by algorithm strength and key length. A common "strong" and high performing cryptographic algorithm today is the Advanced Encryption Standard-Galois Counter Mode (AES-GCM) algorithm. AES-GCM encryption can use keys with 128-, 192- or 256-bit long keys. Longer keys are more secure than shorter keys.[3] Figure 1 shows an overview of the Tape Drive encryption process.

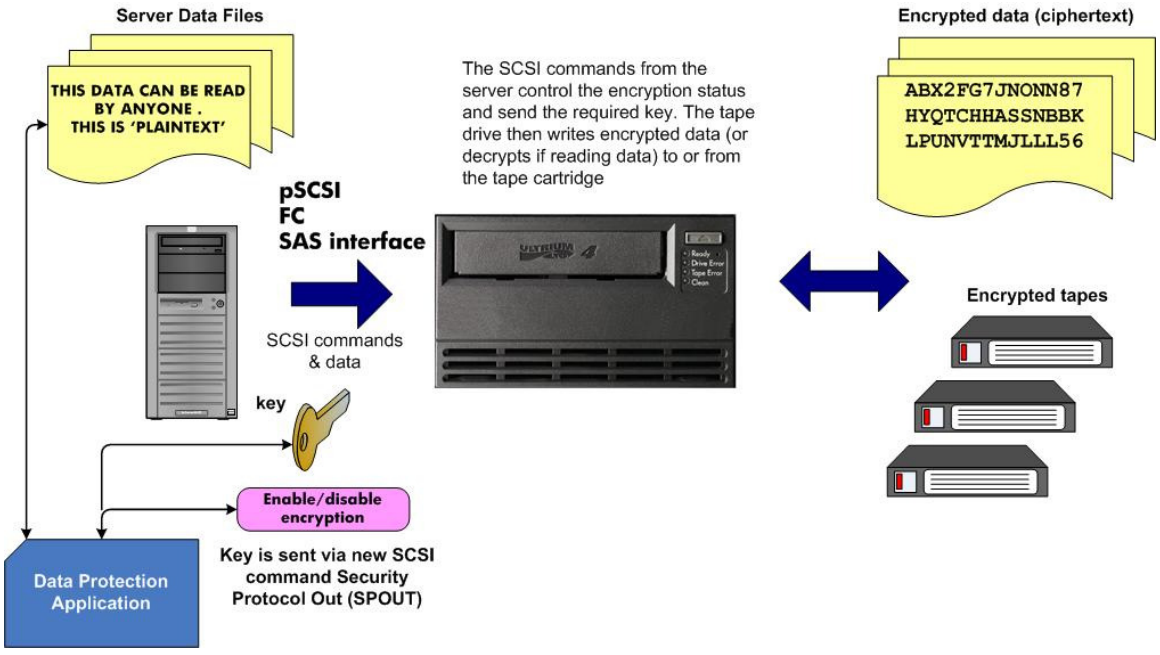

Figure 1 Tape Drive encryption process

[3] CNSS Policy No. 15, available from http://www.cnss.gov/Assets/pdf/cnssp_15_fs.pdf, as of October 27, 2008.

Aside from the length of a key, another important cipher characteristic is that encryption and decryption can be done with the same or different keys.

- **Symmetric or Secret key ciphers** – both encryption and decryption are done with the same key. The advantage of a symmetric key cipher is that encryption and decryption is relatively fast and therefore is typically used to encrypt and decrypt data. This is the technique employed as the LTO-4 encryption method. In Figure 2, Alice wants to send a message to Bob, but does not want it read by third parties. Alice encrypts the message with the one secret key and ensures Bob has the same key to decrypt the data on receipt.

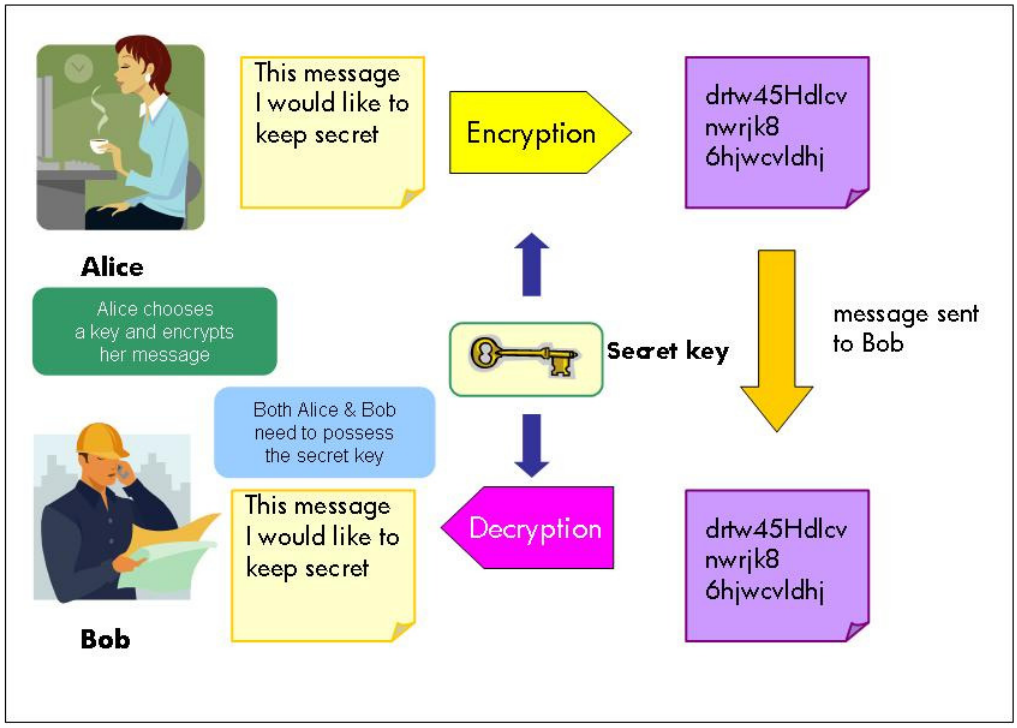

Figure 2 Secret key cipher process

- **Asymmetric or Public key ciphers** – encryption is done with one key and decryption is done with a different but associated key. The advantage of asymmetric key ciphers is that only one party need know the decryption key. Often asymmetric ciphers use a public (encryption) key and a private (decryption) key pair. In Figure 3, Alice wants to send Bob a secret message. Bob generates a public key and private key and keeps his private key to himself, but sends the public key to Alice. Alice encrypts the message using Bob's public key and sends the message to Bob. Bob uses his private key to decrypt the message.

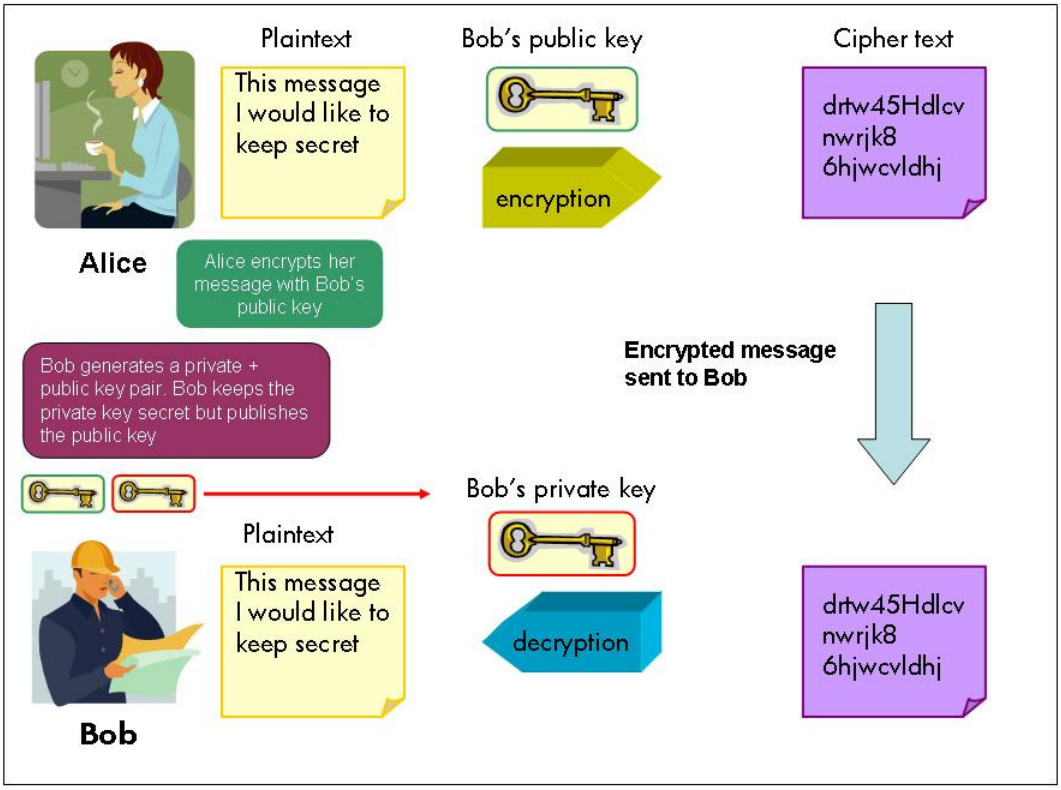

Figure 3 Public key cipher process

Asymmetric key ciphers can also be used in combination with symmetric key ciphers to further secure information. Asymmetric key ciphers are often used to wrap and encrypt symmetric keys for transmission over unsecured networks. The sender, or system, encrypts the symmetric key with the public encryption key and sends it to the receiver, or system, which then uses the private decryption key to decrypt the information. Subsequent data transmissions can then use the much faster symmetric key cipher for encryption/decryption. Asymetric key ciphers in combination with symmetric key ciphers can be used with the LTO-4 encryption process. Symmetric keys can be encrypted or wrapped with asymmetric keys for transmission over IP to the drive, for example, or for symmetric key sharing.

Key management systems

Most data-at-rest solutions, including tape drive encryption, need a key management system that generates, manages, and stores a set of keys accessible within the data center. A key management system can be a separate appliance or more typically is software running on a server in the data center.

Keys generated by the key management system are usually assigned at random to an individual encryption data stream. However, keys may be assigned for the day, or for

any application or any combination thereof. Rotation of multiple different keys would typically be more secure than the use of a single key.

Key management systems store key data in an encrypted "keystore". As such, any keystore data must be backed up to survive in the event of a catastrophic failure. However, keystore data should not be backed up on encrypted tape, as doing so would render all encrypted data including encryption keys unreadable.

For any encryption, key management systems should be configured in high availability clusters to guarantee system availability in the event of hardware failure. Thus, if one key manager fails, another can automatically commence and continuously supply keys without interrupting tape data access.

Key management considerations of tape drive encryption

For tapes, key management systems could store the symmetric key on the cartridge after first securely wrapping it in an asymmetric public key, or alternatively an index or key identification tag for each symmetric key can be written and stored on the encrypted tape(s). This key tag is then used by the key management system to identify the corresponding actual symmetric key, thereby retaining readability of the encrypted tape as well as providing yet another protective layer to the encryption keys.

In situations where a single encrypted tape's information is to be intentionally shared with a "friendly" party, most key management systems would generate a specific export key to enable tape decryption. In fact, this key could be exported and wrapped using an asymmetric key cipher and emailed. Alternatively, it could be burned onto a CD or written to a USB drive and hand delivered to the friendly party. In the latter two cases, the tape and the key should be delivered by two different transport methods to insure that no single interception would garner all the information needed to decrypt the tape.

Because of cost and long standing reliability, tape backups are often relied upon to support disaster recovery operations. Whenever these backup tapes are encrypted, the keystore data also needs to be backed up and safely stored which may include transporting it to the remote site. Compatible hardware, key management systems and possibly applications should also be available at the DR site.

Also, encrypted tapes can be rendered unreadable and effectively erased by deleting the associated encrypting keys from the key management system. Such a key destruction practice can be an important consideration when a tape is no longer in use or is to be recycled for new use by a backup application.

How does LTO-4 tape drive encryption work?

The LTO-4 Tape Drive encryption is specified as part of the LTO-4 open standard format with a 256-bit symmetric key AES-GCM algorithm implemented in tape drive hardware and fully supports the IEEE standard (P1619.1) for tape based encryption and the new

SCSI encryption augmented (T10) command set. The symmetric key is transmitted to the tape transport prior to being used for encrypting data written to or decrypting data read from the media.

The key is not transferred to the tape cartridge and is only retained by the drive during the encryption process. Instead a key identification tag is written and stored on the tape volume. This key identification tag on the tape media provides efficient search access to the necessary information used by the key management system to recall the required encryption key.

Transmission of the keys to the LTO-4 tapes is typically accomplished by using a backup application that supports application managed encryption (AME), by using a tape library that supports library managed encryption (LME), or by using a Key Management Appliance. Most organizations implement LME and tape libraries from IBM, HP, Quantum, Sun, and others support LME tape encryption.

With LME, the tape library has a list of cartridge volume serial numbers that are designated for encryption.

- The backup application requests a mount of a cartridge that is in the library encryption list.
- The library uses the library-to-drive interface to tell the drive to encrypt data on that cartridge.
- The drive requests a symmetric key from the key management software via the libraries IP interface with the key management system and also requests a key tag for the drive to store on the cartridge for subsequent symmetric key identification.

In addition, LME encryption is transparent to the backup application. As such, usually no changes are needed to backup applications. LME can be ideal for environments that have a number of heterogeneous backup applications or servers.

LTO-4 tape libraries can sometimes be partitioned to further support the separation of encrypted from non-encrypted data. Specifically, one or more partitions can be configured to accept only encrypted data whereas the remaining partition(s) only accept non-encrypted data. Some libraries with advanced library management capabilities provide security policy based selection of encryption and specific keys; these can dynamically support a mix of encrypted and non-encrypted cartridges in variable slot locations without needing to use partitions.

Both compression and encryption significantly modify data and can both be performed by an LTO-4 tape drive for the same data on a given tape. In this case, the LTO-4 tape drive first compresses user data and then encrypts it. Thus, the LTO-4 drive can maximize the tape cartridge data capacity and address data security concerns. Also, encrypted data can be added or appended to an LTO-4 encrypted tape cartridge allowing the cartridge capacity to be fully utilized.

Providence Health Customer Case Study

Providence Health and Services is a network of hospitals, health plans, physicians, clinics, home health services, and affiliated health services operating in 5 states throughout the western United States and as such, operates under Health Insurance Portability and Accountability Act (HIPAA) regulations that prescribe how to secure and protect patient medical information.

Providence needed to secure all patient data going offsite, including that on laptops, CDs, USB drives, and tape media. Providence turned to data-at-rest encryption to protect patient data. They now require all laptops, CDs and USB drives to be encrypted when they hold patient data. Additionally, all tapes to be transported offsite are required to be encrypted.

Providence wanted strong reliable encryption and centralized operations. They chose LTO-4's tape drive encryption and library managed encryption in LTO libraries and an Encryption Key Manager (EKM) for its key management system. The EKM runs over a highly available clustered pair of servers and Providence chose to use a rotating set of keys for tape media encryption.

Providence encrypts tapes being transported offsite (Figure 4) and found implementing LTO-4 tape encryption a surprisingly smooth process.

> **"After the tape equipment was installed, it took only 1 to 2 days to implement encryption"**
>
> **Mack Kigada, Data Storage Engineer, Providence Health and Services**

They have now implemented LTO-4 tape with library managed encryption in all 6 of their data centers. In addition, Mr. Kigada, found that LTO-4 encryption had "… virtually no impact on tape drive performance."

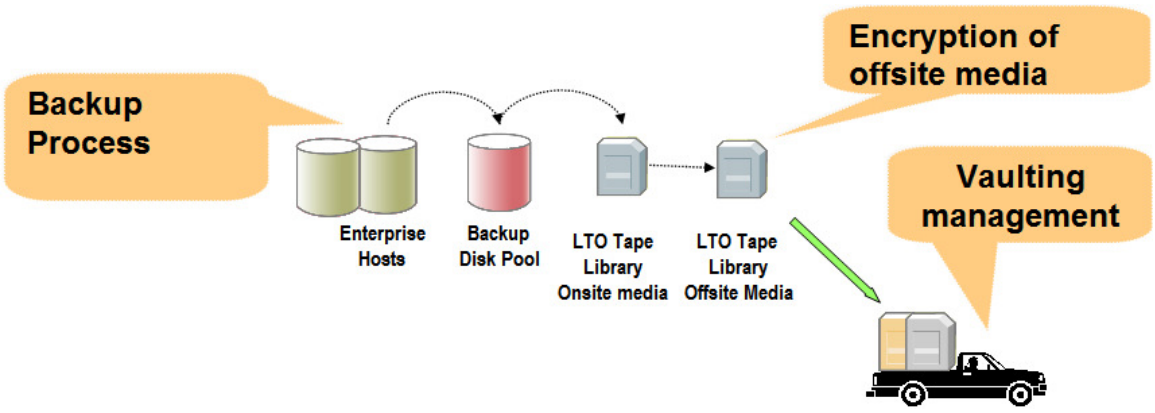

Figure 4 Providence encrypted tape backup process

For offsite tape storage, Providence uses an outside vaulting service. Providence has directed the vaulting service to shred expired unencrypted tapes. Thus, over time all offsite tape data will be encrypted.

Under HIPAA regulations, Providence is subject to periodic HIPAA security audits to insure continuing regulatory compliance. Providence's LTO-4 data encryption combined with other data security measures has passed these continuing security audits without fail.

"The LTO-4 encryption implementation was a well documented process. We installed the key management system ourselves using the supplied documentation. The LTO-4 encryption system has helped us meet our data security policy objectives"

Mack Kigada, Data Storage Engineer, Providence Health and Services

How to implement LTO-4 encryption?

Implementing LTO-4 encryption is a relatively simple, straightforward process. It requires an LTO-4 tape drive, a separate key management system and a backup application or a library that supports AME or LME respectively.

The implementation process is straightforward and requires:

- Generation of initial set of encryption keys by the key management system
- Configuration of key management system and the application or library to provide for two-way transmission of encryption keys and key identification tags
- For AME, enabling of encryption feature or option on the backup application OR for LME, enabling the library to support encryption and then, configuration and identification of encrypting tape drives.

After completion of this process, data can be seamlessly encrypted. However, the key management system has now become critical data center infrastructure and must be treated accordingly.

Typical questions about LTO-4 tape encryption include:

- **Will LTO-4 encryption/decryption affect tape performance?**
  Encryption/decryption performance impact is typically less than 1% and virtually unnoticeable.
- **How many keys do I want to use and how often should they be rotated?**
  Optimal security is realized by providing a key set large enough to assure a different key for every cartridge in an off site shipment and periodic rotation (typically monthly or quarterly) of the key set. This implementation assures there is a different key for every cartridge in a transport set and manages the number of instances of the same key in the offsite storage facility. It additionally makes it possible to use the periodic deletion of keys as a data shredding approach.
- **How long do I need to keep keys?**
  Encryption Keys and their relationship to cartridges must be retained **for as long as tape data needs to be accessed**. This may be as few as a couple of months to as long as 5 to 15 years for tape data archives.

Careful consideration should be given to selection of the Key Management System. Most storage vendors offer key management systems that can be purchased or licensed. Appropriate evaluation criteria should include:

- **Does the key management system support library managed encryption?**
  LME has been the preferred method of encryption implementation since it is typically transparent to the backup application and a variety of servers and backup applications can be supported, normally with no application changes.
- **What key stores are supported?**
  Support for a variety of key store types from existing standards is preferred to allow for flexibility for current and future data security needs.
- **Does the key management system allow for dual cluster support?**
  The key store can typically be backed up for safe keeping but having dual key managers running on separate servers allows for continual operations in the event one becomes unavailable.
- **Can keys be retrieved across a network?**
  This can allow for one key manager server of a clustered pair of servers to be located in another building or across the country for disaster protection and high availability purposes.

Summary

With the proliferation of sensitive data and the ability to store this information indefinitely, the case for encryption is overwhelming. Additionally, current and anticipated future legislation is expected to increase financial penalties and make public notification more onerous for those corporations unwilling or unable to encrypt sensitive data.

Encryption of tape data can safeguard information unwittingly lost or physically stolen. Specifically, data-at-rest encryption securely protects the underlying data of laptops, USB sticks, and tape media. While there are several options for providing data-at-rest tape encryption, the scalability and simplicity of tape drive device encryption would make this option a prime one. LTO-4 vendors provide such encrypting tape drives for the small, mid-range, and enterprise class tape users.

Implementing LTO-4 tape drive encryption is a relatively simple, straightforward process and is well worth the effort. Once implemented there is virtually no tape performance penalty and sensitive data is automatically safeguarded at a fraction of the costs of negative publicity, security breach notification, fines, mandated security measures or civil litigation.

Silverton Consulting, Inc. is a Storage, Strategy & Systems consulting services company, based in the USA offering products and services to the data storage community.

Silverton Consulting
Strategy, Storage & Systems
P.O. Box 492|Broomfield, CO 80038-0492|
720-221-7270|www.SilvertonConsulting.com