

Technology Paper

Self-Encrypting Hard Disk Drives in the Data Center

Data is instantaneously secured the moment the drive leaves the data center.

Introduction

At least 35 U.S. states now have data privacy laws that state if you encrypt data-at-rest, you don't have to report breaches of that data. U.S. Congressional bills have similar provisions. The Payment Card Industry Data Security Standard, which requires rendering sensitive cardholder data unreadable anywhere it is stored, lists strong cryptography as an acceptable method of doing so. For these and other reasons, data center administrators have compelling motivations for encrypting data at rest.

Nearly all drives eventually are decommissioned from use in a data center. When the drives leave the data center, the majority are still operable and their data is still readable. It is not just the drive that is leaving the data center; it's also the data. Even data that has been striped across many drives in a RAID array is vulnerable to data theft. The stripe size in today's arrays is big enough to expose large, intelligible segments of data.

Current drive disposal practices that aim to make the data unreadable are time-consuming and subject to human failure. For this reason, many data centers choose to hire professional disposal services. However, sending a drive offsite for secure disposal puts the drive's data at the same risk as tape data leaving the data center unsecured. Data centers routinely decommission drives, and it only takes one drive to fall through the cracks to possibly cost a company millions of dollars in remedies for the breached data.

The beauty of encryption is that from the moment the drive or system is removed from the data center, whether intentionally or otherwise, the data on the drive is secure. No advance action or thought is required on the part of the data center administrator to secure this data. There is no way for the data to be breached should the drive be mishandled.

Self-Encrypting Hard Disk Drives in the Data Center



Data is instantaneously secured the moment the drive leaves the data center.

Several years ago, before Seagate started working on encryption for the drive, the National Security Agency (NSA) analyzed the problem of data security and determined that the place to perform encryption is in the hard drive. The NSA has spoken publicly at conferences in support of encrypting drives.

However, data center administrators have been reluctant to implement encryption. A major reason for this is fear of losing the key to decrypt their data. Even if they have experience with a good key management system for tape drive encryption, and are convinced that key loss can be prevented, there are concerns about complexity, interoperability, performance and cost. This paper will discuss how self-encrypting hard disk drive technology addresses those concerns.

Technology Overview

This technology consists of three components, described in Figure 1: self-encrypting hard drives, a key management service that stores, manages and serves authentication keys (i.e., passwords), and a storage system that passes these authentication keys to the correct drive.

The self-encrypting drives perform full disk encryption. When a write is performed, clear text enters the drive and, before being written to the disk, is encrypted using an encryption

key embedded within the drive. When a read is performed, the encrypted data on the disk is decrypted before leaving the drive. The black line in Figure 1 denotes the clear text data. The drive requires an authentication key (otherwise known as a password) from an outside source before the drive will unlock for read/write operations.

In addition to its traditional functions, the storage system defines secure volume groups, gets the authentication keys from the key management service, and passes the key to the correct drive. The orange line in Figure 1 denotes this operation. The storage system makes the encryption function transparent to the hosts and applications.

The key management service may include software- or hardware-secure key stores. At the request of the storage system, it will create and assign keys transparently to hosts and applications. The key management service can leverage existing security management policies to define keys and restrict access to keys. Key management includes backup and synchronization, key life-cycle management, auditing, and long-term retention. The key management service can employ existing high-availability and disaster-recovery configurations. The key material can be automatically included within the server backup data and stored offsite.




Figure 1

Self-Encrypting Hard Disk Drives in the Data Center



Data is instantaneously secured the moment the drive leaves the data center.

This technology is designed to be standards-based and to be part of an interoperable solution. The Trusted Computing Group (TCG) Storage Work Group has developed a security communication protocol for self-encrypting drives. Supporting transport commands in T10 and T13 are ratified. IEEE 1619.3 is developing a standard authentication key management protocol. All hard drive vendors and major storage vendors are participating in the Trusted Computing Group. Leading storage system vendors and key management vendors are participating in IEEE 1619.3.

Ultimately, this technology will apply across the entire data center, as shown in Figure 2. Self-encrypting drives may be in storage arrays, on SANs, NAS, and servers, in data centers, branch offices and small businesses. A unified key management service will support the key management requirements for all forms of storage (as well as other security applications).

Proven technology available today is shown highlighted in green in Figure 3. An example is IBM's Key Management for tape encryption. It has been operating with TS1120 encrypting tape

drives for over a year, and it has recently been extended to support encryption on their LTO4 tape drives. Seagate full disk encryption (FDE) drives are currently being used in notebooks, and Seagate has announced a desktop version of those drives.

This paper will now take a closer look at FDE drive technology as it will be used in enterprise drives.

The Hard Drive's Authentication and Encryption Process

It may seem impossible that the encrypting hard drive can provide a secure encryption function without also holding some secret, such as the password or encryption key. Clearly, if it did, any hacker learning a secret would have access to the encrypted data on the drive. In fact, the drive contains no secret that, if discovered, would reveal the encrypted data. A simple description of the unlock process shows how this is true.

The unlock process is the part of the power-on activity that enables access to the encrypted data. The drive expects a credential to be supplied to it, which it verifies as proof that the drive is being accessed by an authorized user.

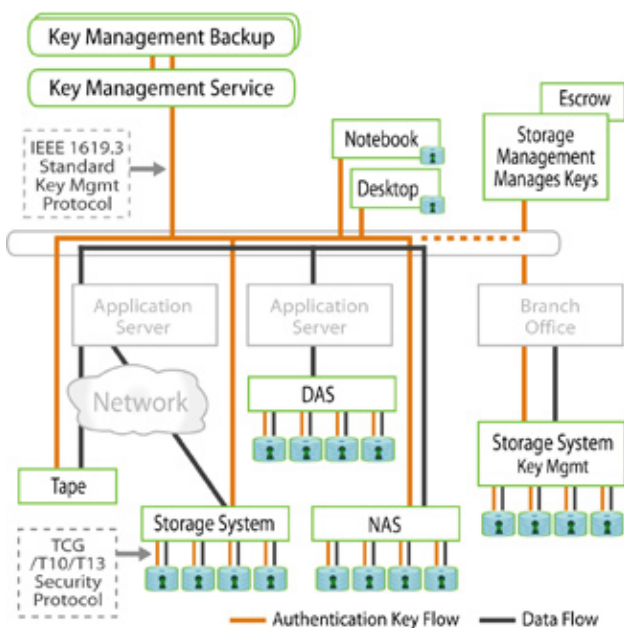


Figure 2




Figure 3

Self-Encrypting Hard Disk Drives in the Data Center

Data is instantaneously secured the moment the drive leaves the data center.



These are the steps of the authentication process with a previously secured drive, as depicted in Figure 4:

1. Authentication

- The storage system gets the authentication key from the key management service and sends it to the correct locked drive.
- The drive hashes the authentication key, and compares that with the hashed authentication key that is stored in a secure area of the disk.
- If the two hashed authentication keys do not match, the authentication process ends, and the drive will not permit reading data from the disk. The drive never sends cipher text from the drive. The drive remains locked.

2. Decrypt the encrypted encryption key

- If the two hashes match, the drive is then unlocked and the drive uses the authentication key it received from the storage system to decrypt a copy of the encryption key, which was previously encrypted with the authentication key. The encrypted encryption key is stored in a secure area of the disk.

This authentication process occurs when the drive is first powered on. After the authentication process is successfully completed, the drive is unlocked until the next time it is powered down. The authentication process does not repeat with each read and write operation. The authentication process does not occur again until the next time the drive is powered up.

3. Clear encryption key encrypts and decrypts the data.

- The clear-text encryption key is then used to encrypt data to be written to the disk, and to decrypt data that is being read from the disk.
- The drive now works in a standard fashion during data transfers with encryption and decryption occurring in the background transparently.

Change of Ownership

This section covers the setup process during initial installation of the drive, repurposing of the drive, and disposal.

The drive generates its own encryption key, using noise from normal internal drive processes as input for random number generation.

When the owner acquires the drive, this embedded encryption key is in clear text form, allowing the end user to use the drive as a normal—non-encrypting—drive. The drive is in an unsecured state. The drive will always encrypt and decrypt all data that it writes to the disk and reads from the disk; however, authentication will not be required, allowing anyone to write and read the clear text data from the disk.

If the owner wishes to use the drive in a secured state, the owner should take steps to lock the drive. To initiate this, the owner should first perform a secure erase, which changes the encryption key. Once the encryption key is changed, any data that had already been written to the disk is unreadable, so normally this operation is performed as soon as the drive is installed. Performing a secure erase prevents a warehouse attack. The owner then establishes a password by entering the SID (proof of ownership) from the drive's external label and sets the password. This password, or authentication key, is used to encrypt the encryption key. The password is also used to lock the drive until the password is given. The drive is now in a secured state. Once the drive is powered down, it will be locked, and when powered up will require authentication before becoming unlocked.

If the owner wishes to repurpose the drive; i.e., change the drive from being in a secured state to an unsecured state so that someone else can




Figure 4

Self-Encrypting Hard Disk Drives in the Data Center

Data is instantaneously secured the moment the drive leaves the data center.



use the drive, the owner issues the secure erase command, which replaces the encryption key and leaves the drive in an unsecured state, just as it was delivered from the factory.

If the drive is believed to be misplaced or stolen, the owner can simply delete all copies of the authentication key, which is equivalent to deleting the encryption key.

Security

When a drive is decommissioned from the data center, the drive is always in danger of falling into the wrong hands.

Figure 5 depicts what potential attackers will have if they obtain this locked drive. Some may assume that because the encryption is performed within the drive, a copy of the encryption key must be kept in the drive. However, as was shown above, the encryption key is not kept in the drive—only an encrypted version of the encryption key is kept in the drive. There are no clear text secrets anywhere on the drive; just a fingerprint (hash) of the password. In designing the drive, Seagate assumes an attacker might have complete knowledge of the drive's design and the location of any secrets held by the drive. Since there is no clue on the drive that could result in deciphering the data, knowing the intricate details of the drive's design and construction in no way benefits hackers. Similarly, breaking one drive provides no secrets that would allow the attacker to break other drives more easily.

Both the data and the encryption key are encrypted using the AES 128 algorithm, the same encryption algorithm approved by the U.S.

government for protecting secret-level classified information.

Generally, cipher text exposure can aid an attack. For instance, if the file system on the drive is a well-known structure, a hacker might be able to take advantage of the fact that certain sectors always contain known values to begin an attack on the encryption. Database structures are also well known. But since the self-encrypting drive does not send cipher text from itself, this type of attack is effectively thwarted. Moreover, the drive can lock itself down after a predetermined number of failed authentication attempts. The drive has protected firmware downloads; an attacker cannot insert modified firmware into the drive.

The self-encrypting drive does not protect against threats within the data center. If an attacker gets access to a server, for example, and the server can access the unlocked drive, then the attacker will be able to read the clear text coming from the drive. This encryption technology does not replace the data center's access controls; rather it is complementary to those controls.

Manageability

The data center administrator doesn't need to escrow the encryption key to maintain data recoverability, because the drive keeps encrypted copies of the encryption key in multiple locations on the drive. If the drive lost all instances of its copies, it is likely caused by drive failure, which makes its data unreadable anyway. Additional encryption keys are automatically added with data redundancy—each time the data is mirrored onto another self-encrypting drive, that drive will have its own set of encrypted encryption keys. Because the data center administrator doesn't need to escrow the encryption key, the key stays protected within the drive.

Separation of authentication and encryption keys provides several management benefits for the user. Because the encryption key itself is encrypted and doesn't leave the drive, the data center administrator doesn't need to change the encryption key periodically; the way a person may change his/her password periodically for security




Figure 5

Self-Encrypting Hard Disk Drives in the Data Center

Data is instantaneously secured the moment the drive leaves the data center.



reasons. That means not having to decrypt and re-encrypt the data, a highly resource-intensive process. The authentication key can be changed as often as desired, without requiring re-encryption. Should a storage administrator leave or a new operator arrive, their rights to access the storage can be incorporated without affecting the encrypted data.

Secure erase involves simply deleting the encryption key or all copies of the authentication key. If the data center administrator wishes to repurpose a drive, the encryption key in the drive can be changed, making the data on the disk unreadable and the drive ready for use by someone else. If the owner wants to return the drive to the manufacturer under warranty, it can be done after performing a secure erase, which lets the manufacturer analyze the drive without exposing the user data. Without encryption, if the owner doesn't want the manufacturer to see the data, they don't even return the drive, losing the benefits of warranty and preventing the manufacturer from learning from the failure to improve future drives.

It's easy to add disk drives with different embedded encryption algorithms to an existing array. The data center can have a mix of encryption algorithms in the same array, because the encryption algorithm is transparent to the system. As drive models change and newer encryption technology gets incorporated into hard drives, they can be intermixed with older drives in storage systems that support encryption without making any changes specific to the new drives' higher level of protection.

Performance

The self-encrypting hard drive uses hardware-based encryption. The encryption engine is in the controller ASIC. There is an encryption engine dedicated to each port on the drive. The encryption engine matches each drive

port's maximum speed. Encryption will not slow the system down. Even better, this scales automatically. As more drives are added, the encryption bandwidth increases commensurately. The data center administrator doesn't have to think about balancing encryption workloads when adding more drives to an array or more arrays to the data center.

Conclusion

Data center administrators have good reason to want to encrypt their data at rest. This technology addresses those reasons and the concerns that have prevented encryption of data at rest in the data center. It is simple: the data is secure the moment a drive is removed from the system. A drive may be compromised, but it will not expose the data.

The fear of losing the authentication key and being unable to decrypt one's own data is addressed with proven key management capabilities, operational in the largest financial businesses. Keys can be securely backed up, replicated and mirrored in disaster-recovery centers. Self-encrypting hard drives also address questions of complexity, interoperability, performance and cost. The technology is designed to be standards-based for optimal manageability and interoperability, and all major hard drive manufacturers are participating in the standards activity. Also, standards drive volume and competition, which drives cost. This technology is designed to be integrated into standard products. The encryption is transparent to normal storage management and end users. Performance scales linearly, automatically.

Encryption in the drive provides superior performance, manageability and security. For that reason, many analysts, system manufacturers and government agencies like the NSA are suggesting it should be done there. The bottom line: This is a significant leap forward to improve security and management in the world's data centers.

AMERICAS Seagate Technology LLC 920 Disc Drive, Scotts Valley, California 95066, United States, 831-438-6550
ASIA/PACIFIC Seagate Technology International Ltd. 7000 Ang Mo Kio Avenue 5, Singapore 569877, 65-6485-3888
EUROPE, MIDDLE EAST AND AFRICA Seagate Technology SAS 130-136, rue de Sully, 92773, Boulogne-Billancourt Cedex, France 33 1-4186 10 00

Copyright © 2007 Seagate Technology LLC. All rights reserved. Printed in USA. Seagate, Seagate Technology and the Wave logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. All other trademarks or registered trademarks are the property of their respective owners. Seagate reserves the right to change, without notice, product offerings or specifications. Seagate does not guarantee the accuracy of the information contained on this document. The information provided on this document does not constitute legal advice and should not be relied upon as such. You should consult with your own legal counsel for such advice. Additionally, the information contained on this document is current as of the date indicated herein and as such this document may not reflect changes in the law following this date. Publication Number: TP583.1-0711US, November 2007