Education

ABCs of Data Encryption for Storage

Eric A. Hibbard, CISSP, CISA, ISSAP, ISSMP, ISSEP, SCSE
Hitachi Data Systems

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract

ABCs of Data Encryption for Storage

Public disclosures of data "indiscretions" have become regular enough and embarrassing enough that many organizations are exploring encryption options to simply stay out of the headlines. Those who have ventured into this space quickly realize that there is no "magic crypto fairy dust" that will make the problems go completely away. However, with careful planning and judicious use of the right technologies, organizations can eliminate many of their exposures.

This session focuses on the efforts required at the storage layer to create a successful encryption strategy. Major uses along with factors to consider are presented for protecting storage management, data in-flight, and data at-rest. The session provides expanded coverage on encrypting data at-rest, including key management and a step-by-step approach.

Encryption Basics
(What You Need to Know)

Check out SNIA Tutorial:

Cryptography Deciphered

A Few Definitions

- Plaintext – Original information (intelligible) that is used as input to an encryption algorithm (cipher).
- Ciphertext – The encrypted (unintelligible) output from an encryption algorithm.
- Encryption – The conversion of plaintext to encrypted text (ciphertext) with the intent that it only be accessible to authorized users who have the appropriate decryption key.
- Cipher – A mathematical algorithm for performing encryption (and the reverse, *decryption*).
- Key – A piece of auxiliary information used by a cipher during the encryption operation.

General Categories of Encryption Algorithms

- Symmetric-key Ciphers (Secret-key Cryptography)
  - Uses the same key to encrypt and decrypt the data
  - Two types: block ciphers & stream ciphers
  - Block ciphers commonly used for storage
  - Generally much less computationally intensive than asymmetric-key ciphers
- Asymmetric-key Ciphers (Public Key Cryptography)
  - Use a pair of keys with a mathematical association that allows any data encrypted by one key to be decrypted only by the other.
  - Often used for authentication & digital signatures rather than encrypting data
- Hashing Algorithms
  - Does not encrypt data, but provides a one-way (non-reversible) transformation used to store data securely as well as to verify data integrity
  - Does not require the use of keys
  - The size of the value output by the hashing process is fixed (SHA-1 is 20 bytes)

Encryption Introduction

- Goals of Encryption
  - Make data unintelligible to unauthorized readers
  - Make it extremely difficult to decipher data when attacked
- Factors to consider:
  - Strength of encryption (algorithm, key size)
  - Quality of encryption (sufficiently reviewed by experts; implementations subjected to accreditation)
  - Speed of encryption
  - Management of the persistent encryption keys
  - Randomness (use of random number generator)

Why Use Encryption?

- Legal obligations that require privacy
- Regulatory requirements that include data confidentiality compliance elements (in the form of encryption measures)
- Due care mandates the use of confidentiality measures to protect valuable data assets
- National security is dependent on the secrecy of certain data

Fundamentals of Encryption

- Encryption does not, in and of itself, provide anything other than data confidentiality.
- Most effective when used with authentication, access control, and integrity measures
- Confidentiality of encrypted data is dependent on keeping the encryption key secret, rather than keeping the inner workings of the algorithm or cipher secret.
- Encryption can introduce additional complexities due to export/import restrictions (not limited to vendors)

SNIA Position on Encrypting "Sensitive" Data

- Externalized data (data leaving your control)
  - Data stored on removable media (like backup tapes), which potentially leaves the control of the organization, must be encrypted while at-rest
  - Data stored in third-party datacenters must be encrypted both in-flight and at-rest within these "untrusted" datacenters
  - Data transferred between "trusted" datacenters (controlled by the organization) must be encrypted
- Encrypting Data At-rest – A measure of last resort
  - Use extreme care when encrypting primary data
  - Long-term key management is a critical element

Sensitive Data = Data that require special protection due to legal, regulatory, statutory, and/or competitive requirements.

Encryption and Audit Compliance

- Information Systems Auditors are required to understand:
  - The encryption "basics" and their application
  - Risk Assessment in use of Encryption
  - Encryption Legislation/Regulation
- Information criteria most relevant to an encryption technology audit:
  - Primary: Effectiveness, confidentiality, integrity, availability and compliance
  - Secondary: Efficiency and reliability
- Suggested Procedures (ISACA Guidance to Auditors):
  - Organizational Management – Written Procedures/Policies
  - Change control over the cryptographic system including key management
  - Design criteria of a cryptographic system
  - Digital Signature – Controls and proper use
  - Validity conditions of a cryptographic algorithm

Source: Information Systems Audit and Control Association (ISACA) Evaluation Procedure P9

Data Leakage Example

Original

Tux © Larry Ewing (lewing@isc.tamu.edu)

Encrypted using ECB mode

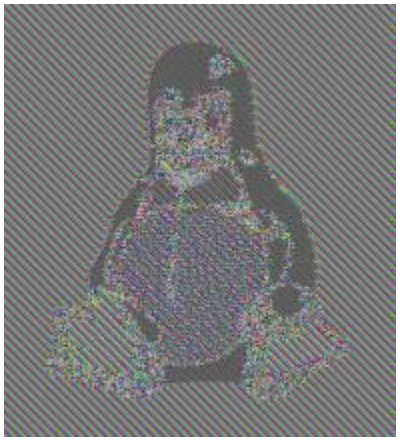

Encrypted using CBC or CTR modes

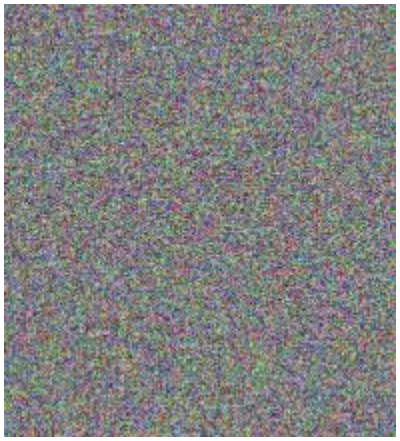

http://en.wikipedia.org/wiki/Block_cipher_modes_of_operation

ECB = electronic codebook mode
CBC = cipher-block chaining mode
CTR = counter mode

Storage Layer Encryption

In-flight versus At-rest

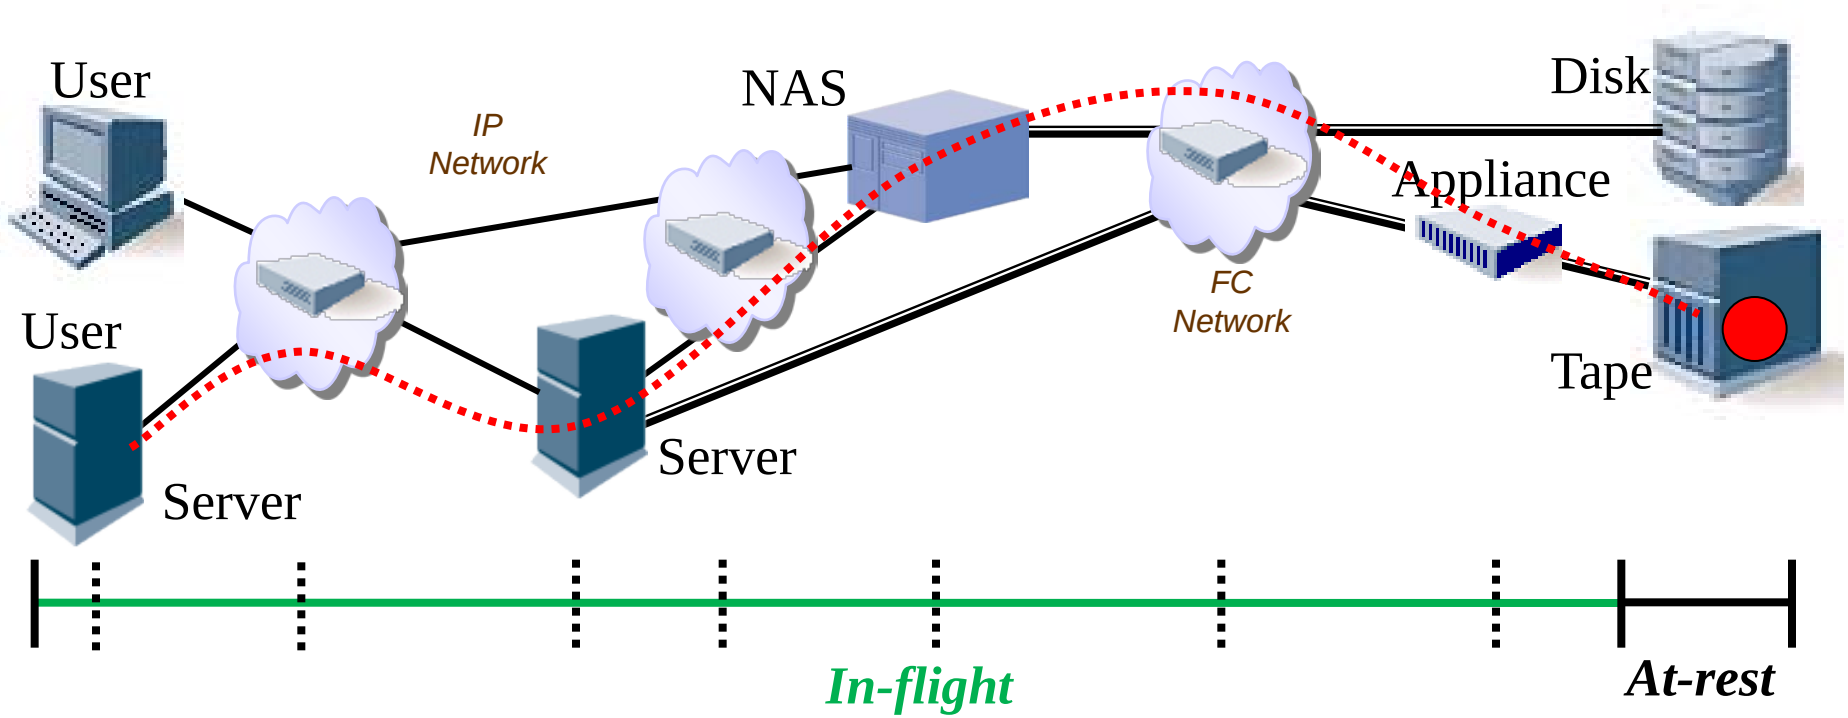

In-flight:
- Two end points (communication)
- Interoperability – network layers
- Data is transitory (temporary)

At-rest:
- Interoperability – media interchangeability
- Data is persistent on media

Very Different Threats and Threat Agents

Applying Encryption to Storage

- Protecting Data In-flight
  - Data Access
    - Block-level, IP protocols (IPsec for iSCSI, iFCP, FCIP)
    - Block-level, FC protocols (FC-SP ESP_Header)
    - File-level, IP protocols (IPsec for NFS & SMB/CIFS; SSL/TLS for WebDAV)
  - Management (IPsec, TLS, SSH)
- Protecting Data At-rest
  - Block-level storage (FC, iSCSI, FCoE)
  - File-level storage (NFS, CIFS, pNFS)

Protecting Data At-rest

- Used to protect against unauthorized disclosure of sensitive data while they are resident in storage media (disk, tape, optical)
- May provide in-flight protection, depending on the point of encryption
- Factors to consider:
  - Sensitivity/importance of the data (directly impacts the solution)
  - Long-term key management; escrow services
  - Retention and destruction
  - Access requirements (e.g., multi-user access to a single file)
  - Placement of the point of encryption
  - Impacts on disaster recovery/business continuity implementations
  - Proof of encryption (e.g., when the tape falls off the truck)

Prerequisites to Encrypting Data

- Data Classification:
  - Different types of sensitive/important data require different types of protection
  - Role of compliance to meet regulatory and legal obligations
- Understand Where the Data Reside:
  - Inventory data assets; map to storage/media
  - Analyze data flows (intermediate copies); cradle-to-grave (DR/BC & archives are often overlooked)
- Understand How the Data Are Accessed/Used:
  - Degree of shared data access (e.g., group access to a file)
  - Official versus unofficial versions/copies of data (downloaded copies to laptops)
- Understand Potential Operational Issues:
  - Potential impacts to data availability (e.g., backups & DR/BC)
  - Potential impacts to performance
  - Potential impacts to scalability
  - Proof of encryption

Where to Apply Encryption

- Security Perspective: Encrypt as close to the source as possible.
- Points of Encryption (one, some, or all):
  - **Application-level** – under the control of specific application or database; finest granularity of control and maximum insight into the data (type, users, sensitivity)
  - **Filesystem-level** – under the control of the OS or OS-level application; control at file-level with insights into the users
  - **Network-level** – under the control of a network-based system
    - File-based (NAS) – control at the share/filesystem-level (possibly file-level) with moderate insights into the users
    - Block-based – control at the logical volume level with limited insights in the "community of users"
  - **Device-level** – under the control of the end-device; control at the logical volume level with limited insights in the "community of users"

Points of Encryption

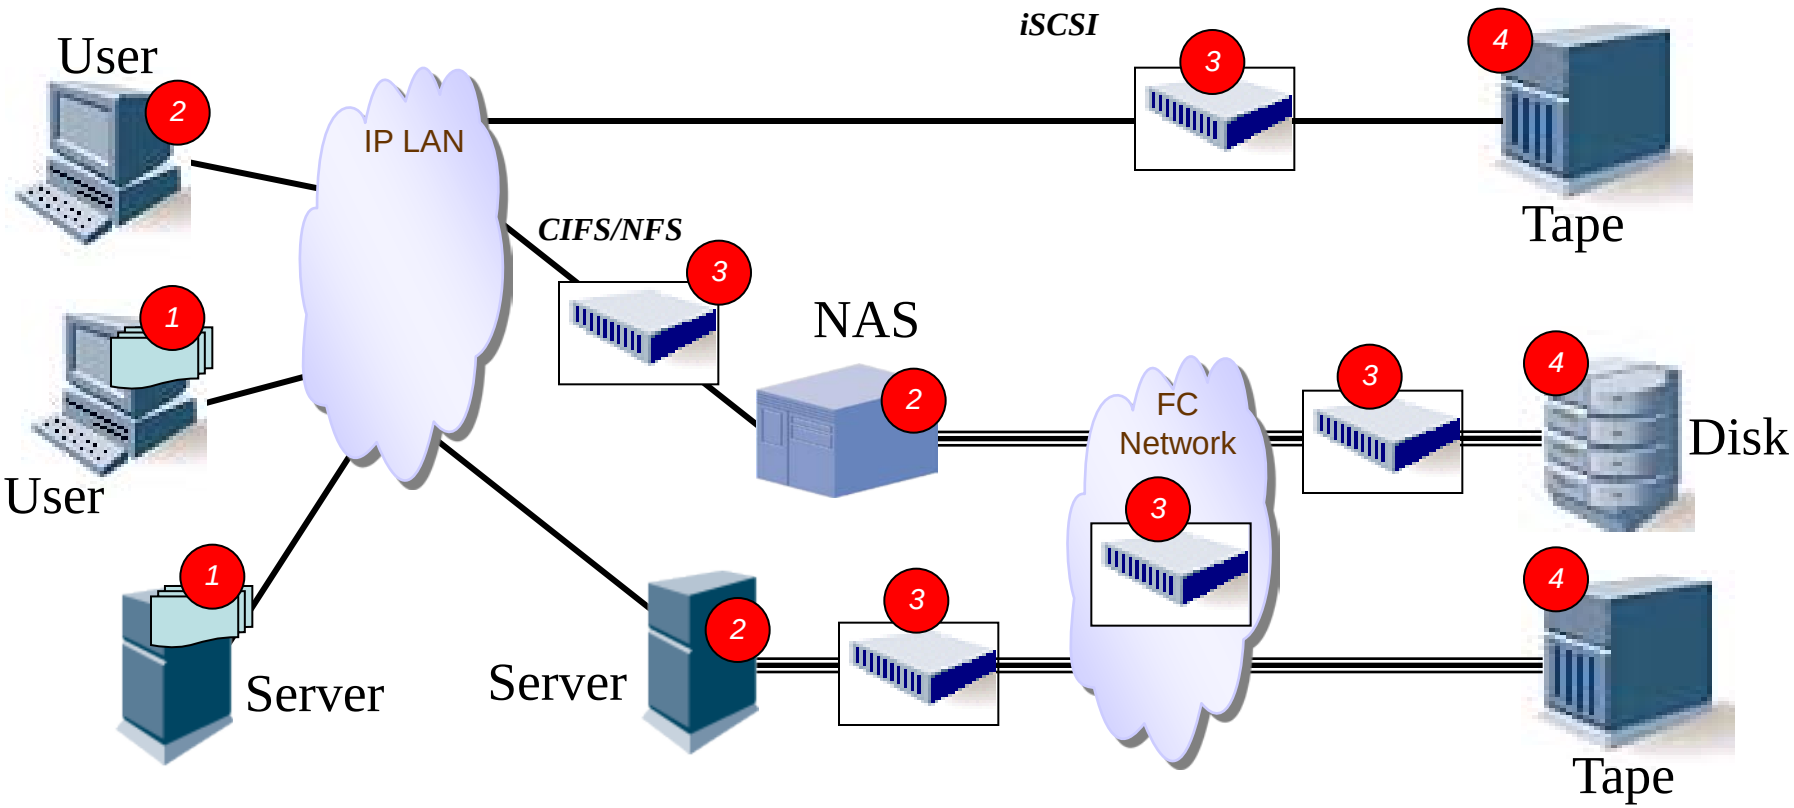

Potential Areas of Impact (1)

- User: User sees a change in the interface, process, and/or storage mechanism
- Availability: The degree to which the overall availability of the system/solution will be restricted, diminished, or eliminated.
- Infrastructure: Networking, systems, and storage infrastructure (e.g., moving LUNs) must be changed
- Performance/Throughput: Negative impact compared to existing (low=10%, moderate=20%, significant=35%, extreme=50%+)
- Scalability: The degree to which the overall scalability of the existing system will be restricted, diminished, or eliminated.

Potential Areas of Impact (2)

- In-flight Confidentiality: Characterization of in-flight confidentiality protection from the user system/application to the storage device
- Business Continuity/Disaster Recovery (BC/DR): The degree to which the overall BC/DR will be restricted, diminished, or eliminated.
- Proof of encryption: Characterization of the proof of encryption aspects (e.g., functionality, integration into existing infrastructure, evidence)
- Environmentals: Characterization of the environmental aspects (e.g., power, cooling, space)

Comparison of Impacts

IMPACT	APPLICATION	FILESYSTEM	NETWORK	DEVICE
User	Low	Low-Moderate	None	None
Availability	Can be significant	Can be significant	Low-Moderate (Redundancy)	Low-Moderate
Infrastructure	Can be significant	Can be significant	Low-Moderate	Low
Performance/ Throughput	Can be severe	Can be significant	Low	Low-Moderate
Scalability	Can be significant	Can be significant	Can be moderate	Minimal
In-flight Confidentiality	Excellent	Low-Moderate (NAS); Excellent (Host)	Low-Moderate	None
BC/DR	Can be extremely complicated	Can be complicated	Can be extremely complicated	Can be extremely complicated
Proof of Encryption	Can be complicated	Relatively easy	Low-Moderate	Can be complicated
Environmentals	Low-Moderate	Low-Moderate	Can be significant	Low

Managing the Keys
(Encrypting Data At-rest)

Key Management 101

- Definition: The *activities* involving the handling of cryptographic keys and other related security parameters (e.g., initialization vectors and passwords) during the entire life cycle of the keys, including their generation, storage, establishment, entry and output, and destruction. [NIST SP 800-57 Part 1]
- Appropriate and successful key management is critical to the secure use of *every* crypto system without exception.
- Considered the most difficult aspect of cryptography because of the human element (involves system policy, user training, organizational and departmental interactions, coordination between end users, etc.)

Key Management Operations

- Generation – Creation of fully random keys
- Distribution – Keys have to be adequately protected (encrypted) when they are transmitted over networks
- Storage – Keys are encrypted wherever they are stored on some form of media; decryption of one key should not expose others in the process
- Recovery – Ability to restore (e.g., from backup or escrow service) a key that has been lost or corrupted
- Destruction – Ability to permanently destroy an encryption key, rendering the encrypted data unusable

Key Management Guidance (1)

- A cryptographic key should be used for only one purpose.
- Use appropriate key wrapping (encryption of symmetric keys) when keys are transmitted and/or stored.
- Limit the amount of time a key is in plaintext form and prevent humans from viewing plaintext keys.
- Use a cryptographic integrity check on the keys.
- NIST SP 800-57 Part 1 identifies minimum symmetric security levels, defined in bits of strength (not key size)
  - 80 bits of security until 2010 (128-bit AES and 1024-bit RSA)
  - 112 bits of security through 2030 (3DES, 128-AES and 2048-bit RSA)
  - 128 bits of security beyond 2030 (128-AES and 3072-bit RSA)

Key Management Guidance (2)

- A symmetric data encryption key should be used no more than 2 years to protect (encrypt) data
- When data are retained as ciphertext for extended periods and/or the key may have been compromised, the data should be re-keyed (decrypted and then encrypted, using a new key)
- Have a compromise recovery plan in the event of a key compromise.
- Destroy keys, rather than expire, as soon as they are no longer needed.

SNIA Encryption Checklist

SNIA Security Whitepaper, *Encryption of Data At-rest – Step-by-step Checklist*, http://www.snia.org/ssif/documents

A Step-by-Step Approach

- Defines a set of tasks to effectively implement at-rest encryption
  - Defines a process, not a single activity
  - Not all sub-steps will be needed, but they all merit consideration
- Identifies important sources of encryption requirements
  - Payment Card Industry (PCI) Data Security Standard (DSS)
  - Information Systems Audit and Control Association (ISACA)
  - Federal Financial Institutions Examination Council (FFIEC)

Summary of the Steps (1)

1. Understand confidentiality drivers
   - Identify all relevant regulatory and other obligations
   - Identify all relevant legal obligations
   - Identify all relevant executive management concerns
   - Review organizational policies
   - Review organizational IS/IT strategic plans
2. Classify the data assets
   - Assume it may not be able to encrypt everything
   - Identify the organizational classifications of the high-value & sensitive data
   - Determine the organization's confidentiality categories & priorities
   - Focus on a small number of coarse classifications

Summary of the Steps (2)

3. Inventory data assets
   - For each confidentiality category determine
     - The hosts & applications that process the data
     - The data owners, custodians, stakeholders, and business units
     - The devices that store the data and their geographic locations
     - The networks which are used to transport the data
   - Perform a risk assessment
4. Perform data flow analysis
   - Identify temporary and permanent storage locations
   - Consider data protection schemes (backups, CDP, replication, etc.)
   - Determine the impact of data reduction schemes (compression and deduplication)
   - Consider the role of mobile devices

Summary of the Steps (3)

5. Determine the points-of-encryption

- Encrypt as close to the source as possible to maximize protections
- Determine the granularity needed for the encryption
- Consider both in-flight and at-rest requirements
- Determine the risks to be mitigated by the encryption solution
- Identify the preferred point-of-encryption for each category

6. Design encryption solution

- Develop and document the organization's encryption strategy and architecture/framework as well as requirements
- Determine the key management and proof-of-encryption needs
- Factor in business continuity and/or disaster recovery measures
- Set the selection criteria and document potential impacts

Summary of the Steps (4)

Education SNIA

7. Begin data re-alignment, if required

- Data may need to be migrated or re-aligned to take full advantage of the expected encryption solution
- Identify specific data to be relocated and develop an action plan to re-align this data
- Adjust data protection schemes and related CDP, DLP, compression & deduplication processes
- Begin the data re-alignment efforts

8. Implement solution

- Determine the approach to field the encryption solution
- Select technology and acquire/develop the components
- Deploy and integrate the encryption & key management technology
- Complete end-to-end testing and prepare roll-back plan

Summary of the Steps (5)

9. Activate encryption

- Complete informal or formal management accreditation of the encryption solution
- Turn on the actual encryption capabilities
- Run point tests to prove that the data can be processed & recovered
- Complete final data re-alignment activities

Done!

Final Thoughts

Additional Considerations

- Encryption as a method for media sanitization (a.k.a., crypto shredding/erasing)
  - Encryption key is destroyed, rendering data unusable
  - Requires careful use of symmetric encryption keys (e.g., one key per file/directory/file system/LUN)
  - ALL copies of the key, including escrows and backups, must be destroyed
  - Not permitted for national security data
- Selection of cipher and mode should be carefully considered to address threats and possible data leakage

Remember…

- Determining the primary driver (compliance versus data security) for encryption is critical
- Classification of the organizational data can significantly improve the effectiveness of most encryption solutions
- The "need" for encryption, combined with insufficient money, often results in impacts to business processing
- Key management complexities are almost always overlooked, but they are critical success factors for the encryption solution
- Interoperability is not guaranteed, so attention to detail is important
- For all that encryption offers, it does not come for free

Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: tracksecurity@snia.org

Many thanks to the following individuals for their contributions to this tutorial.

- SNIA Education Committee

Eric A. Hibbard, CISSP, CISA
Larry Hofer, CISSP, PE
Roger Cummings
Tim Smith
Richard Austin, CISSP
Andrew Nielsen, CISSP, CISA
Ray Kaplan, CISSP
Gianna DaGiau

SNIA Security TWG

SNIA Security Organizations

- SNIA Security Technical Work Group (TWG)
  - Focus: Requirements, architectures, interfaces, practices, technology, educational materials, and terminology for storage networking.
  - http://www.snia.org/tech_activities/workgroups/security/
- Storage Security Industry Forum (SSIF)
  - Focus: Educational materials, customer needs, whitepapers, and best practices for storage security.
  - http://www.snia.org/ssif