



Sanitization

March 2015

Abstract: *The ISO/IEC 27040:2015 (Information technology - Security techniques - Storage security) standard provides detailed technical guidance on controls and methods for securing storage systems and ecosystems. This whitepaper provides an overview of the sanitization guidance in the standard as applied to disk-type media and provides guidance to organizations in developing a sanitization program to meet their particular needs.*

USAGE

The SNIA hereby grants permission for individuals to use this document for personal use only, and for corporations and other business entities to use this document for internal use only (including internal copying, distribution, and display) provided that:

1. Any text, diagram, chart, table or definition reproduced shall be reproduced in its entirety with no alteration, and,
2. Any document, printed or electronic, in which material from this document (or any portion hereof) is reproduced shall acknowledge the SNIA copyright on that material, and shall credit the SNIA for granting permission for its reuse.

Other than as explicitly provided above, you may not make any commercial use of this document, sell any or this entire document, or distribute this document to third parties. All rights not explicitly granted are expressly reserved to SNIA.

Permission to use this document for purposes other than those enumerated above may be requested by e-mailing tcmd@snia.org. Please include the identity of the requesting individual and/or company and a brief description of the purpose, nature, and scope of the requested use.

All code fragments, scripts, data tables, and sample code in this SNIA document are made available under the following license:

BSD 3-Clause Software License

Copyright (c) 2014, The Storage Networking Industry Association.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

* Neither the name of The Storage Networking Industry Association (SNIA) nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

DISCLAIMER

The information contained in this publication is subject to change without notice. The SNIA makes no warranty of any kind with regard to this specification, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The SNIA shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this specification.

Suggestions for revisions should be directed to <http://www.snia.org/feedback/>.

Copyright © 2015 SNIA. All rights reserved. All other trademarks or registered trademarks are the property of their respective owners.

Revision History

Revision	Date	Sections	Originator:	Comments
<i>V1.0 R0</i>	<i>3/4/2015</i>	All	Richard Austin	Initial Draft
<i>V1.0 R1</i>	<i>3/5/2015</i>	All	Richard Austin	Final Draft

Suggestion for changes or modifications to this document should be submitted at <http://www.snia.org/feedback/>.

Foreword

This is one of a series of whitepapers prepared by the SNIA Security Technical Working Group to provide an introduction and overview of important topics in ISO/IEC 27040:2015, *Information technology – Security techniques – Storage security*. While not intended to replace the standard, they provide additional explanations and guidance beyond that found in the actual standard.

Executive Summary

Sanitization addresses a risk as old as reusable storage devices themselves: disclosure of information via reuse or disposal of storage media. Though originally as easy as assuring that magnetic media was degaussed before reuse or disposal, advances in technology and new storage media types have increased the complexity of the sanitization process while society as a whole is becoming increasingly intolerant of data stewards who put their information at risk.

This document provides an overview of the sanitization guidance provided in ISO/IEC 27040 and will assist organizations in developing a sanitization program that meets their needs. As sanitization decisions may have implications for retention, discovery holds and other regulatory or legal matters, the users of this document are advised to consult with competent legal counsel before implementing its guidance. Nothing in this document should be construed as legal advice or opinion.

1 Introduction

The purpose of sanitization is, not too surprisingly, to sanitize storage media. “Sanitize” means rendering the data on the media infeasible to recover for a given level of effort.¹ The “level of effort” is an important concept as it will motivate decisions regarding the sanitization process and its implementation. As will be covered later, the more valuable the information, the more effort an adversary might be willing to invest in recovering it and the more careful an organization must be in its choice of a sanitization method.

Sanitization is required by four basic types of concerns:

1. In many jurisdictions, loss of control of media containing protected information is being regarded as equivalent to disclosing that information. This may mean that breach notifications must be performed and liability for regulatory or other legal penalties may be incurred.
2. Developing privacy requirements (such as the European Union “Right to be forgotten”) may require that data not be just “deleted” but also sanitized to assure it is resistant to recovery.
3. It may be a requirement. Defense contractors and other agencies that deal in national security information must comply with regulations governing sanitization.
4. Limiting the scope of information subject to electronic discovery – organizations generally want to limit the amount of information that is potentially subject to

¹ ISO/IEC 27040, p. 5

electronic discovery requests. This may be accomplished by assuring that all data at end-of-life (including retention periods) and not subject to a current duty to preserve is sanitized.²

Whether, and how, these concerns apply to your organization is determined through consultation with competent legal counsel. As this is a rapidly developing area of law, this consultation should be regular and ongoing rather than a one-time occurrence.

This whitepaper will primarily focus on disk-type media not because other storage types are less important or ISO/IEC 27040 lacks guidance for other media types but because readers will likely be more familiar with disk-type media. A sound grasp of sanitization concepts in the familiar disk-media context will prepare readers to successfully apply ISO/IEC 27040's guidance for other storage types.

1.1 History

As noted earlier, sanitization is as old as reusable storage media. An early (1991) discussion of the problem revolved around “remanence” – “the residual information that remains on storage media after erasure.”³ Remanence recognized that operations intended to delete data (such as a file system delete operation) might leave recoverable parts of the data behind that could be recovered by an interested party. Though the Trusted Computer System Evaluation Criteria (TCSEC) has been superseded, this document is the source for much of the terminology for sanitization methods.

An influential study of how data could be reconstructed from its remanence and effective methods for accomplishing sanitization was conducted by Peter Gutmann in 1993.⁴ Gutmann also described sophisticated laboratory-style data recovery methods such as magnetic force microscopy and developed the Gutmann algorithm for sanitizing media that provided assurance against them. This method of overwriting multiple times with varying data patterns became the recommended guidance for many sanitization products and was incorporated in mandatory guidance such as the 1995 version of the NISP Operating Manual or DoD 5220.22-M. However, as noted in an epilogue to his original paper, later technological developments have modified the need for multiple overwrites in many cases.

² As noted earlier, sanitization goes beyond simple deletion and includes resistance to some level of effort.

³ Trusted Computer System Evaluation Criteria *A Guide to Understanding Data Remanence in Automated Information Systems*, the “forest green” book. <http://fas.org/irp/nsa/rainbow/tg025-2.htm> Accessed 30 January, 2015

⁴ Peter Gutmann, *Secure Deletion of Data from Magnetic and Solid-State Memory*. Sixth USENIX Security Symposium. https://www.cs.auckland.ac.nz/~pgut001/pubs/secure_del.html with epilogue accessed 30 January, 2015.

Though sanitization was a concern for the national security and military establishments, the private sector and individuals had much less understanding of its requirements. This was graphically changed when Simpson Garfinkel purchased a large number of used disk drives and did a study of the data that was easily retrieved from them.⁵ Garfinkel retrieved a trove of confidential data such as medical and financial records.

A major modern study of disk sanitization was conducted at the Center for Magnetic Recording Research (CMRR) at UCSD in 2007.⁶ This study provides an excellent overview of the various requirements for sanitization, effective methods for sanitizing disk media and is notable for its coverage of the new device commands in modern drives that provide sanitization capability and also debunks a number of exaggerated claims regarding the ability to recover data from sanitized disk media. It is also notable for identifying the limitations of degaussing as a sanitization method due to the increasing use of solid-state memory for storage devices (and as performance-enhancing caches in magnetic storage devices).

NIST's Special Publication 800-88, *Guidelines for Media Sanitization* (newly revised in December of 2014) provides detailed guidance for sanitizing many forms of media ranging from disks to paper. Though its guidance is binding only on US federal agencies, it is an excellent resource for private sector organizations.

ISO/IEC 27040 provides extensive guidance for storage security in general and includes sanitization as a major topic. It's normative (i.e., describing what should be done) Annex A: *Media Sanitization* provides detailed guidance on how to properly sanitize many types of storage based on the risk profile of the organization.

1.2 Implement sanitization as a process

Sanitization is not a one-off activity but rather should be implemented as a normal part of the information management process. A process implementation should also include appropriate proof of sanitation to both satisfy audit requirements and justify exclusion from electronic discovery requests (e.g., if user workstation disks are automatically sanitized within 30 days of termination of employment, no legal duty to preserve or produce information from Jane Doe's workstation would be created by a lawsuit filed 60 days after her termination⁷).

⁵ Simpson Garfinkel and Abhi Shelat, *Remembrance of Data Passed: A Study of Disk Sanitization Practices*. IEEE Security & Privacy, January/February 2003. pp. 17-27.

⁶ Gordon Hughes and Tom Coughlin, *Tutorial on Disk Drive Sanitization*, 2007.
<http://cmrr.ucsd.edu/people/Hughes/documents/DataSanitizationTutorial.pdf> Accessed January 6, 2015

⁷ This is an overly simplified example – if the organization reasonably expected the litigation to be filed earlier than 30 days after Jane's termination, a duty to preserve might already exist in advance of the lawsuit. This is the type of legal issue that mandates regular consultation with legal counsel on sanitization decisions and processes.

1.2.1 Legal Considerations

As noted earlier, the sanitization process must comply with both retention periods (whether mandated by law, regulation or policy) and the duty to preserve. If a given type of information (e.g., executive email) must be retained for some period, then the sanitization process must exclude it until the retention period has expired.

Duty to preserve deals with the responsibility of a party to a litigation to preserve any information relevant to that litigation. Though there is a general good-faith acceptance that normal sanitization processes may destroy potentially relevant information, there is a requirement to suspend sanitization when the litigation is filed or even when a reasonable expectation of litigation arises. This implies that there should be a means for the organization's legal staff to request suspension of sanitization activities and for that suspension to be implemented in a timely fashion.

This is an active and developing area of law and regular consultation with legal counsel is required.

2 Methodology

2.1 Considerations in choosing a methodology

Choosing a sanitization methodology requires balancing the value of the information against the capabilities and intent of an adversary⁸ as well as deciding on the desired final disposition of the media.

Valuing information can be a difficult process as most private sector organizations are not subject to the strict classification protocols of the national security establishment. Private sector organizations tend to think in terms of contractual, legal and regulatory mandates that require certain levels of protection for specific types of information (e.g., PCI DSS for payment card information, HIPAA for protected health information and privacy requirements for personally identifiable information). Information on new products, significant inventions, etc., also have a direct value that must be protected. Even when information is valued, the identification of storage media holding specific types of information can be difficult. In many cases, this state of affairs, unfortunately, requires sanitizing storage to a level appropriate to the most sensitive information it *might* contain.

⁸ "Adversary" is used generically to represent anyone desiring illicit access to valuable information whether they are cybercriminals, operatives of a state-sponsored intelligence operation, cyber-activists, etc.

Adversary intent and capability also vary. Cybercriminals seek to monetize pilfered payment card information either directly by making fraudulent purchases or indirectly by selling it to those who will make those purchases. Nation states may use pilfered intellectual property for competitive advantage. Capability can range from use of simple undelete utilities to modern forensic tools to techniques available only in a well-funded and equipped laboratory.

The importance of the final disposition of media revolves around whether the sanitized media will remain under organizational control or be sold, recycled⁹, returned to the vendor, etc.

Though there are many variables and their values may be nebulous, the choice of sanitization methodology revolves around two simple questions:

1. Will your organization retain control of the media after it is sanitized?
2. What level of technical capability do potential adversaries have?

2.1.1 Retaining Control

If the media will be reused within your organization, then its exposure to outside entities is limited. Your risk here is more on the lines of the curious employee or contractor rather than a cybercriminal gang with significant technical expertise and capability. In this case a less strenuous sanitization method would be appropriate.

On the other hand, if the sanitized media is leaving your control, you have no idea who may attempt to retrieve information. Used hard drives from a financial institution would likely attract the interest of cybercriminals hoping to recover payment card or other valuable financial information. Similarly, used media from a defense contractor would be of interest to nation states engaged in espionage.

A sometimes overlooked case involves media that is returned to a vendor under a warranty claim or service contract. This case is complicated by the possibility that the device may not be functional at the time it is returned.

Media leaving an organization's control should be subjected to a more stringent sanitization regimen.

2.1.2 Adversary capability

Attempts to recover data from sanitized media break down into two basic categories: logical methods and physical methods such as sophisticated laboratory attacks.

⁹ Storage devices contain many useful materials (such as rare earth elements) that can be recovered by a suitably equipped recycling center.

Logical methods range from using a file undelete utility to common digital forensic tools. These tools are readily available and any adversary may be assumed to have access to them.

Physical methods are open ended and depend solely on the technical expertise of the adversary and the effort they are willing to expend. As noted earlier, Gutmann described the use of Magnetic Force Microscopy to retrieve data from magnetic disks (though this technique no longer works with modern drives). A modern example would be using laser drills to access the internal circuitry of a logic chip and a logic analyzer to decode its operation.¹⁰ These potential attacks are particularly relevant to solid state storage where wear-leveling may leave significant information extant at hardware levels below the storage controller.

When adversary capability is high and the information is valuable enough to justify use of those advanced capabilities, the media should be subjected to the most stringent sanitization regimens.

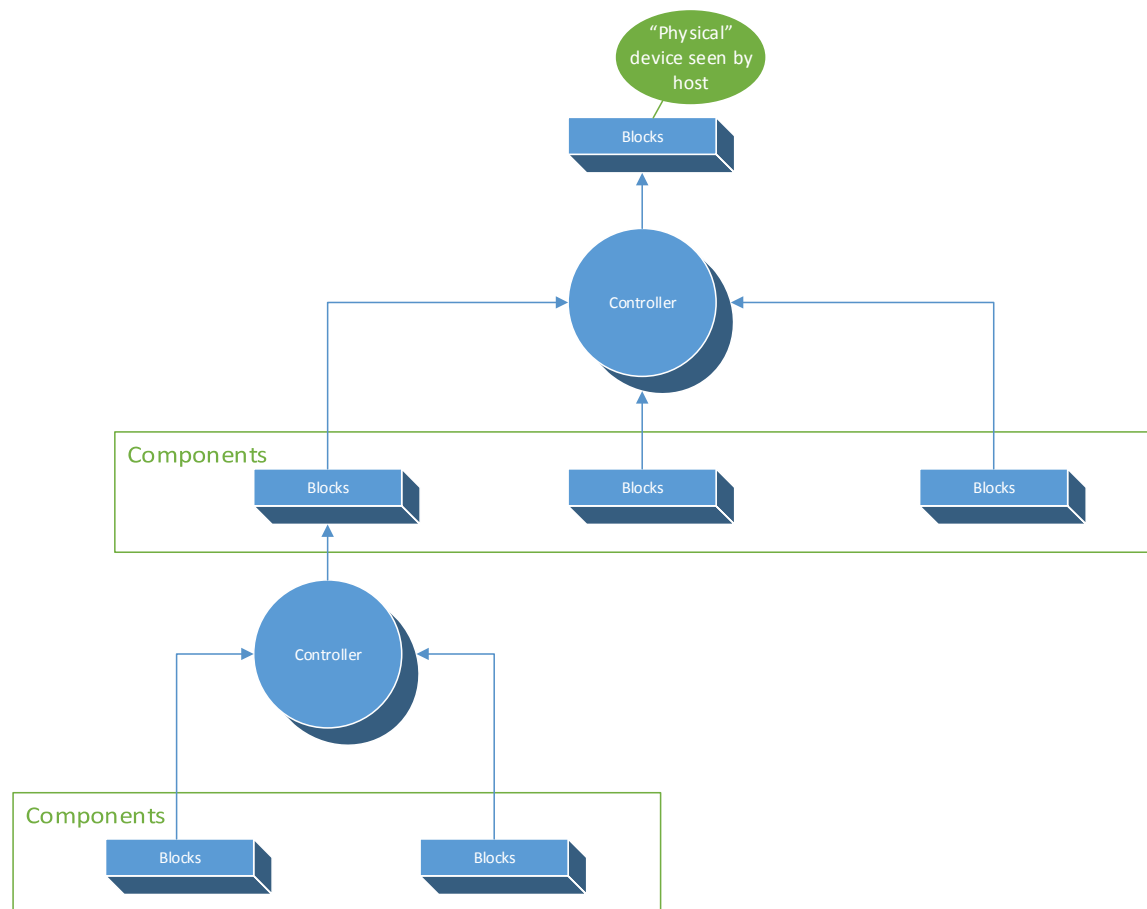


Figure 1: Levels of Virtualization in Storage Ecosystems

¹⁰ For a good overview of such laboratory attacks, see Chapter 16 of Ross Anderson's *Security Engineering: A Guide to Building Dependable Distributed Systems* (2nd ed). 2008.

2.2 Logical versus Media-Aligned Sanitization

Modern storage ecosystems can be complex environments that implement several layers of virtualization in presenting what appears to be a disk device to a host. Looking at Figure 1, at the top are a collection of blocks presented to a host as a storage LUN. Below that is a controller which combines chunks of storage blocks (possibly from multiple independent devices) to create that LUN seen by the host. Perhaps one of those devices (the one on the left) is a solid state drive so the physical “device” is actually a virtual creation made up of a collection of blocks that are managed by a subsidiary controller (for purposes of wear-leveling, etc.).

Sanitization can be applied at any of these levels and ISO/IEC 27040 uses the terminology logical and media-aligned to differentiate between them. If one were to sanitize the virtual drive (the LUN), that would be an example of *logical* sanitization. If one were to identify the subsidiary devices used by the controller in creating the virtual drive and sanitize those, this would be *media-aligned sanitization*.

Logical sanitization is more commonly done because it is much easier to identify and access the virtual devices dedicated to a particular host than to parse out which chunks of storage on which physical drives contribute space to the higher level entity.

2.3 Terminology

Sanitization methodologies fall into three general classes: clear, purge and destruct. Clear is the least stringent while destruct, not too surprisingly, is the most stringent. Clear and purge leave the media in a reusable state while destruct does not.

2.3.1 Clear

ISO/IEC 27040 defines clear as “sanitize using logical techniques on data in all user-addressable storage locations for protection against simple non-invasive data recovery techniques using the same interface available to the user”. The clear operation is most often implemented using an overwriting process which fills every user-addressable block of the device with non-sensitive data (such as binary 0’s). As noted earlier, on modern magnetic disks, a single overwrite pass is sufficient.

Clear is appropriate for media that must withstand logical attacks but will not be subject to physical attacks. This is basically due to the limitations inherent on only overwriting the user-addressable portions of the media as it excludes portions of the media that may be reserved and sections that may have been deleted from the usable space by sparing operations.

A clear operation is also not appropriate for solid state drives as technological features such as wear leveling may transparently (at the host level) move logically accessible data to different physical locations while leaving the previous location's content intact.

Though it may seem a bit confusing, organizations may choose to use the device specific commands (or even cryptographic erase) normally considered as appropriate for purge operations (see below) to perform a clear. This is basically due to convenience as these techniques typically execute much faster than an overwriting process. To help avoid confusion, remember that these techniques can be used for clear operations as a *matter of convenience* while purge operations rely on them because of their increased *effectiveness*.

2.3.2 Purge

ISO/IEC 27040 defines purge as “sanitize using physical techniques that make recovery infeasible using state of the art laboratory techniques, *but which preserves the storage media in a potentially reusable state*”. A significant characteristic of purge operations is that they sanitize both user-addressable and non-user-addressable storage locations. This may be accomplished using specialized device commands or degaussing (for magnetic media).

Purge operations are appropriate when the media will be passing out of organizational control (recycled, donated, resold) as the organization cannot be sure of what level of effort a potential adversary may devote to retrieving information potentially remaining on the media.

Cryptographic erase

As media has become larger and as new technologies such as solid-state storage have come into widespread use, organizations have sought to minimize the time required to perform sanitization even on devices having sanitization commands. The increasing use of encryption for data at rest has provided a very fast technique for cryptographic erasure. Cryptographic erasure basically involves destroying the encryption key for the data and thus forcing an adversary to conduct an attack against the cryptologic implementation in order to gain access to the sanitized data. Another advantage of cryptographic erase is its high granularity. For example, it is theoretically possible to cryptographically erase a single field in a database by encrypting it under a random key that is immediately sanitized.

Cryptographic erase may also be the only effective technique for sanitizing certain types of media (such as flash-based solid-state storage).

Though it sounds deceptively easy to implement and very attractive because of its speed, there are a number of critical requirements that must be met:

Encryption must be applied before any data is written to the drive – this requirement assures that there is no data on the media in clear-text form.

High-pedigree encryption is required. This requires that the cryptographic algorithms themselves and their implementation must be reliable. This requirement is usually met by using only a well-vetted cryptographic implementation such as those that have met FIPS140-2 (or equivalent) certification. This is to assure that there are no weaknesses in the implementation that would make it easier for an adversary to access the encrypted information.

Effective key management is required. In order to conduct cryptographic erasure, **all** copies of the relevant key must be sanitized (including those in escrow or a centralized key management solution).¹¹

Proof of encryption is required. In order for cryptographic erase to be accepted as a sanitization method, it must be reliably documented that the data was encrypted appropriately in the first place.

Meeting these requirements assures that a potential adversary must either mount a brute-force attack on the encrypted data or find a weakness in the algorithm or its implementation (highly unlikely in a well-vetted implementation).

2.3.3 Destruct

ISO/IEC 27040 defines destruct as “sanitize using physical techniques that make recovery infeasible using state of the art laboratory techniques and *results in the subsequent inability to use the media for storage.*” Destruct is commonly implemented by means such as incineration, shredding, etc.

Though it sounds deceptively simple (and possibly even fun), care must be taken in selecting an appropriate method to match the risk of data exposure. Field expedient means such as firing a bullet through a disk might preclude its being put back in a server to read the data but the damaged platters might still provide useful data to a well-funded and technically adept adversary.

Whatever method is chosen should also consider environmental impacts accruing to what is done with the media (e.g., landfill waste), possible toxic fumes produced during incineration, etc.

¹¹ For further information, please refer to the whitepaper in this series: Encryption/Key Management

2.3.4 Summary

Selecting a sanitization method, like many other decisions in information security, is based on risk. Table 1 provides a concise summary of where each method is appropriate.

Method	Media End State	Adversarial Risk Level
Clear	Reusable	Low-Medium
Purge	Potentially Reusable	High
Destruct	Not Reusable	Very High

Table 1 Method vs. Adversarial Risk

Adversarial risk reflects media exposure (under/out of organization control), content information value and adversary capability.

Once the appropriate method has been identified, *Annex A* of ISO/IEC 27040 provides guidance in choosing the appropriate technique for a specific type of media.

3 Conclusion

Sanitization is a critical component of information lifecycle management to assure that information is not inadvertently disclosed when storage is reused, recycled or discarded. ISO/IEC 27040 provides extensive guidance in selecting methods and techniques for many types of storage based on the risk profile of the organization.

4 Acknowledgments

4.1 About the Author

Richard Austin has worked in the IT industry for over 35 years in positions ranging from software developer to security architect. He is currently employed as a security architect in the Technology Office of HP Cyber Security. He is a senior member of both the IEEE and ACM and also holds the CISSP certification. In addition to active participation in SNIA's Security Technical Working Group, he participates in international standardization activities through INCITS/CS1.

4.2 Reviewers and Contributors

The Security TWG wishes to thank the following for their contributions to this whitepaper:

Eric Hibbard, CISSP, CISA

Hitachi Data Systems

Walt Hubis

Hubis Technical Associates

Dr. Alan Yoder

Huawei Technologies Co. Ltd.

Wayne Adams

EMC Corporation

5 For More Information

Additional information on SNIA security activities, including the Security TWG, can be found at <http://www.snia.org/security>.

Suggestion for revision should be directed to <http://www.snia.org/feedback/>.

The ISO/IEC 27040 standard can be purchased at <http://www.iso.org>.